
MarkLogic Server

MarkLogic Connector for Hadoop Developer's Guide

MarkLogic 5
October, 2011

Last Revised: 5.0-3, March, 2012

MarkLogic Connector for Hadoop Developer's Guide

| | | |
|-------|---|----|
| 1.0 | Introduction to MarkLogic Connector for Hadoop | 6 |
| 1.1 | Terms and Definitions | 6 |
| 1.2 | Overview | 8 |
| 1.2.1 | Job Building Tools Provided by the Connector | 9 |
| 1.2.2 | Input and Output Selection Features | 9 |
| 1.2.3 | MarkLogic Server Access via XDBC App Server | 10 |
| 1.3 | MarkLogic-Specific Key and Value Types | 10 |
| 1.4 | Deploying the Connector with a MarkLogic Server Cluster | 11 |
| 1.4.1 | Relationship of MarkLogic Server to a Hadoop Cluster | 11 |
| 1.4.2 | Jobs Use In-Forest Evaluation | 12 |
| 1.4.3 | Cluster-wide XDBC Configuration Requirements | 12 |
| 1.5 | Making a Secure Connection to MarkLogic Server with SSL | 13 |
| 2.0 | Getting Started with the MarkLogic Connector for Hadoop | 15 |
| 2.1 | Requirements | 15 |
| 2.1.1 | Required Software | 15 |
| 2.1.2 | Security Requirements for MapReduce Jobs | 16 |
| 2.2 | Installing Apache Hadoop and the MarkLogic Connector for Hadoop | 16 |
| 2.2.1 | Installing Apache Hadoop | 16 |
| 2.2.2 | Installing the MarkLogic Connector for Hadoop | 18 |
| 2.3 | Configuring Hadoop to Use the Connector | 18 |
| 2.4 | Running the HelloWorld Sample Application | 19 |
| 2.4.1 | Creating the App Server and Database | 19 |
| 2.4.2 | Loading the Input Data | 20 |
| 2.4.3 | Configuring the Job | 21 |
| 2.4.4 | Running the Job | 22 |
| 2.5 | Making the Connector Available Across a Hadoop Cluster | 23 |
| 2.6 | Organization of the Connector Distribution | 23 |
| 3.0 | Apache Hadoop MapReduce Concepts | 25 |
| 3.1 | MapReduce Overview | 25 |
| 3.2 | Example: Calculating Word Occurences | 27 |
| 3.3 | Understanding the MapReduce Job Life Cycle | 28 |
| 3.3.1 | Job Client | 28 |
| 3.3.2 | Job Tracker | 29 |
| 3.3.3 | Task Tracker | 29 |
| 3.3.4 | Map Task | 30 |
| 3.3.5 | Reduce Task | 30 |
| 3.4 | How Hadoop Partitions Map Input Data | 31 |
| 3.5 | Configuring a MapReduce Job | 32 |
| 3.5.1 | Configuration Basics | 32 |
| 3.5.2 | Setting Properties in a Configuration File | 33 |

| | | |
|------------|--|-----------|
| 3.5.3 | Setting Properties Using the Hadoop API | 33 |
| 3.5.4 | Setting Properties on the Command Line | 34 |
| 3.5.5 | Configuring a Map-Only Job | 34 |
| 3.6 | Running a MapReduce Job | 34 |
| 3.7 | Viewing Job Status and Logs | 34 |
| 4.0 | Using MarkLogic Server for Input | 36 |
| 4.1 | Basic Steps | 36 |
| 4.1.1 | Identifying the Input MarkLogic Server Instance | 36 |
| 4.1.2 | Specifying the Input Mode | 37 |
| 4.1.3 | Specifying the Input Key and Value Types | 38 |
| 4.1.4 | Defining the Map Function | 39 |
| 4.2 | Basic Input Mode | 40 |
| 4.2.1 | Creating Input Splits | 40 |
| 4.2.2 | Using a Lexicon to Generate Key-Value Pairs | 41 |
| 4.2.2.1 | Implement a Lexicon Function Wrapper Subclass | 42 |
| 4.2.2.2 | Override Lexicon Function Parameter Wrapper Methods | 42 |
| 4.2.2.3 | Choose an InputFormat | 44 |
| 4.2.2.4 | Configure the Job | 45 |
| 4.2.2.5 | De-duplication of Results Might Be Required | 46 |
| 4.2.3 | Using XPath to Generate Key-Value Pairs | 46 |
| 4.2.4 | Example: Counting Href Links | 48 |
| 4.3 | Advanced Input Mode | 50 |
| 4.3.1 | Creating Input Splits | 50 |
| 4.3.1.1 | Overview | 50 |
| 4.3.1.2 | Creating a Split Query with <code>hadoop:get-splits</code> | 51 |
| 4.3.2 | Creating Input Key-Value Pairs | 52 |
| 4.3.3 | Optimizing Your Input Query | 53 |
| 4.3.4 | Example: Counting Hrefs Using Advanced Mode | 54 |
| 4.4 | Using <code>KeyValueInputFormat</code> and <code>ValueInputFormat</code> | 56 |
| 4.4.1 | Overview | 57 |
| 4.4.2 | Job Configuration | 57 |
| 4.4.3 | Supported Type Transformations | 58 |
| 4.4.4 | Example: Using <code>KeyValueInputFormat</code> | 58 |
| 4.5 | Configuring a Map-Only Job | 60 |
| 4.6 | Input Configuration Properties | 60 |
| 4.7 | InputFormat Subclasses | 62 |
| 4.8 | Lexicon Function Subclasses | 64 |
| 5.0 | Using MarkLogic Server for Output | 66 |
| 5.1 | Basic Steps | 66 |
| 5.1.1 | Identifying the Output MarkLogic Server Instance | 66 |
| 5.1.2 | Configuring the Output Key and Value Types | 67 |
| 5.1.3 | Defining the Reduce Function | 68 |
| 5.1.4 | Disabling Speculative Execution | 69 |

| | | |
|---------|--|-----|
| 5.1.5 | Example: Storing MapReduce Results as Nodes | 70 |
| 5.2 | Creating a Custom Output Query with KeyValueOutputFormat | 72 |
| 5.2.1 | Output Query Requirements | 73 |
| 5.2.2 | Job Configuration | 73 |
| 5.2.3 | Supported Type Transformations | 75 |
| 5.3 | Controlling Transaction Boundaries | 77 |
| 5.4 | Performance Considerations for ContentOutputFormat | 77 |
| 5.4.1 | Time vs. Space: Configuring Batch and Transaction Size | 77 |
| 5.4.2 | Time vs. Correctness: Using Direct Forest Updates | 78 |
| 5.5 | Output Configuration Properties | 79 |
| 5.6 | OutputFormat Subclasses | 83 |
| 6.0 | Troubleshooting and Debugging | 85 |
| 6.1 | Enabling Debug Level Logging | 85 |
| 6.2 | Solutions to Common Problems | 85 |
| 6.2.1 | Configuration File Not Found | 86 |
| 6.2.2 | XDBC App Server Not Reachable | 86 |
| 6.2.3 | Authorization Failure | 86 |
| 7.0 | Using the Sample Applications | 87 |
| 7.1 | Set Up for All Samples | 87 |
| 7.1.1 | Install Required Software | 87 |
| 7.1.1.1 | Multi-host Configuration Considerations | 88 |
| 7.1.2 | Install the Sample Configuration Files | 88 |
| 7.1.3 | Modify the Sample Configuration Files | 88 |
| 7.2 | Additional Sample Data Setup | 90 |
| 7.2.1 | Creating the Database | 90 |
| 7.2.2 | Creating the XDBC App Server | 91 |
| 7.2.3 | Loading the Data | 91 |
| 7.3 | Interacting with HDFS | 92 |
| 7.3.1 | Initializing HDFS | 93 |
| 7.3.2 | Accessing Results Saved to HDFS | 94 |
| 7.3.3 | Placing Content in HDFS to Use as Input | 94 |
| 7.4 | Sample Applications | 95 |
| 7.4.1 | HelloWorld | 97 |
| 7.4.2 | LinkCountInDoc | 98 |
| 7.4.3 | LinkCountInProperty | 98 |
| 7.4.4 | LinkCountValue | 99 |
| 7.4.5 | LinkCount | 100 |
| 7.4.6 | LinkCountCooccurrences | 100 |
| 7.4.7 | RevisionGrouper | 101 |
| 7.4.8 | BinaryReader | 102 |
| 7.4.9 | ContentReader | 103 |
| 7.4.10 | ContentLoader | 104 |
| 7.4.11 | ZipContentLoader | 105 |

| | | |
|-----|-------------------------|-----|
| 8.0 | Technical Support | 107 |
| | Product Notices | 108 |
| | COPYRIGHT | 108 |
| | TRADEMARK NOTICE | 115 |

1.0 Introduction to MarkLogic Connector for Hadoop

The MarkLogic Server Hadoop MapReduce Connector provides an interface for using a MarkLogic Server instance as a MapReduce input source and/or a MapReduce output destination.

This chapter is an overview of the MarkLogic Connector for Hadoop, covering:

- [Terms and Definitions](#)
- [Overview](#)
- [MarkLogic-Specific Key and Value Types](#)
- [Deploying the Connector with a MarkLogic Server Cluster](#)
- [Making a Secure Connection to MarkLogic Server with SSL](#)

If you are not already familiar with Apache Hadoop MapReduce, see “Apache Hadoop MapReduce Concepts” on page 25.

For installation instructions and an example of configuring and running a job, see “Getting Started with the MarkLogic Connector for Hadoop” on page 15.

1.1 Terms and Definitions

You should be familiar with the following terms and definitions before using the Hadoop MapReduce Connector.

| Term | Definition |
|--------------------------------|--|
| <i>Apache Hadoop MapReduce</i> | An Apache Software Foundation software framework for reliable, scalable, distributed parallel processing of large data sets across multiple hosts. The Hadoop core framework includes a shared file system (HDFS), a set of common utilities to support distributed processing, and an implementation of the MapReduce programming model. See “Apache Hadoop MapReduce Concepts” on page 25. |
| <i>job</i> | The top level unit of work for a MapReduce system. A job consists of an input data set, a MapReduce program, and configuration properties. Hadoop splits a job into map and reduce tasks which run across a Hadoop cluster. A Job Tracker node in the Hadoop cluster manages MapReduce job requests. |
| <i>task</i> | An independent subcomponent of a job, performing either map or reduce processing. A task processes a subset of data on a single node of a Hadoop cluster. |

| Term | Definition |
|--------------------------|---|
| <i>map task</i> | A task which contributes to the map step of a job. A map task transforms the data in an <i>input split</i> into a set of output key-value pairs which can be further processed by a reduce task. A map task has no dependence on or awareness of other map tasks in the same job, so all the map tasks can run in parallel. |
| <i>reduce task</i> | A task which contributes to the reduce step of a job. A reduce task takes the results of the map tasks as input, produces a set of final result key-value pairs, and stores these results in a database or file system. A reduce task has no dependence on or awareness of other reduce tasks in the same job, so all the reduce tasks can run in parallel. |
| <i>mapper</i> | Programatically, a subclass of <code>org.apache.hadoop.mapreduce.Mapper</code> . The mapper transforms map input key-value pairs into map output key-value pairs which can be consumed by reduce tasks. An input pair can map to zero, one, or many output pairs. |
| <i>reducer</i> | Programmatically, a subclass of <code>org.apache.hadoop.mapreduce.Reducer</code> . The reducer aggregates map output into final results during the reduce step of a job. The value portion of an input key-value pair for reduce is a list of all values sharing the same key. One input key-value pair can generate zero, one, or many output pairs. |
| <i>input source</i> | A database, file system, or other system that provides input to a job. For example, a MarkLogic Server instance or HDFS can be used as an input source. |
| <i>input split</i> | The subset of the input data set assigned to a map task for processing. Split generation is controlled by the <code>InputFormat</code> subclass and configuration properties of a job. See “How Hadoop Partitions Map Input Data” on page 31. |
| <i>input split query</i> | When using MarkLogic Server as an input source, the query that determines which content to include in each split. By default, the split query is built in. In advanced input mode, the split query is part of the job configuration. |
| <i>input query</i> | When using MarkLogic Server as an input source, the query that generates input key-value pairs from the fragments/records in the input split. |

| Term | Definition |
|----------------|---|
| InputFormat | <p>The abstract superclass, <code>org.apache.hadoop.mapreduce.InputFormat</code>, of classes through which input splits and input key-value pairs are created for map tasks.</p> <p>The Apache Hadoop MapReduce API includes <code>InputFormat</code> subclasses for using HDFS as an input source. The MarkLogic Connector for Hadoop API provides <code>InputFormat</code> subclasses for using MarkLogic Server as an input source; see “InputFormat Subclasses” on page 62.</p> |
| OutputFormat | <p>The abstract superclass, <code>org.apache.hadoop.mapreduce.OutputFormat</code>, of classes that store output key-value pairs during the reduce phase.</p> <p>The Apache Hadoop MapReduce API includes <code>OutputFormat</code> subclasses for using HDFS for output. The MarkLogic Connector for Hadoop API provides <code>OutputFormat</code> subclasses for using a MarkLogic Server database as an output destination; see “OutputFormat Subclasses” on page 83.</p> |
| <i>HDFS</i> | <p>The Hadoop Distributed File System, which can be used as an input source or an output destination in jobs. HDFS is the default source and destination for Apache Hadoop MapReduce jobs.</p> |
| <i>shuffle</i> | <p>The process of sorting all map output values with the same key into a single <code>(key, value-list)</code> reduce input key-value pair. The shuffle happens between map and reduce. Portions of the shuffle can be performed by map tasks and portions by reduce tasks.</p> |

1.2 Overview

This section provides a high level overview of the features of The MarkLogic Connector for Hadoop. If you are not already familiar with Apache Hadoop MapReduce, you should first read “Apache Hadoop MapReduce Concepts” on page 25.

Topics covered in this section:

- [Job Building Tools Provided by the Connector](#)
- [Input and Output Selection Features](#)
- [MarkLogic Server Access via XDBC App Server](#)

1.2.1 Job Building Tools Provided by the Connector

The MarkLogic Connector for Hadoop manages sessions with MarkLogic Server and builds and executes queries for fetching data from and storing data in MarkLogic Server. You only need to configure the job and provide map and reduce functions to perform the desired analysis.

The MarkLogic Connector for Hadoop API provides tools for building MapReduce jobs that use MarkLogic Server, such as the following:

- `InputFormat` subclasses for retrieving data from MarkLogic Server and supplying it to the map function as documents, nodes, and user-defined types. See “InputFormat Subclasses” on page 62.
- `OutputFormat` subclasses for saving data to MarkLogic Server as documents, nodes and properties. See “OutputFormat Subclasses” on page 83.
- Classes supporting key and value types specific to MarkLogic Server content, such as nodes and documents. See “MarkLogic-Specific Key and Value Types” on page 10.
- Job configuration properties specific to MarkLogic Server, including properties for selecting input content, controlling input splits, and specifying output destination and document quality. See “Input Configuration Properties” on page 60 and “Output Configuration Properties” on page 79.

1.2.2 Input and Output Selection Features

Using MarkLogic Server for input is independent of using it for output. For example, you can use a MarkLogic Server database for input and save your results to HDFS, or you can use HDFS for input and save your results in a MarkLogic Server database. You can also use MarkLogic Server for both input and output.

The MarkLogic Connector for Hadoop supports two input modes, basic and advanced, through the `mapreduce.marklogic.input.mode` configuration property. The default mode is basic. In basic input mode, the connector handles all aspects of split creation, and your job configuration specifies which content in a split is transformed into input key-value pairs.

In advanced input mode, you control both the split creation and the content selection by writing an input split query and an input query. For details, see “Using MarkLogic Server for Input” on page 36. Basic mode provides the best performance.

When using MarkLogic Server for input, input data can come from either database content (documents) or from a lexicon.

MapReduce results can be stored in MarkLogic Server as documents, nodes, and properties. See “Using MarkLogic Server for Output” on page 66.

1.2.3 MarkLogic Server Access via XDBC App Server

The MarkLogic Connector for Hadoop interacts with MarkLogic Server through an XDBC App Server. When using MarkLogic Server for both input and output, the input server instance and output server instance can be different. The connector API includes configuration properties for identifying the server instances.

The configured MarkLogic Server instance acts as an initial point of contact for the job, but the MarkLogic Connector for Hadoop spreads the query load across all nodes in the MarkLogic Server cluster that host a forest of the target database. The job communicates directly with each node, rather than bottlenecking on the single MarkLogic Server instance configured into the job. For details, see “Deploying the Connector with a MarkLogic Server Cluster” on page 11.

The MarkLogic Connector for Hadoop creates and manages the XDBC sessions for your application using XCC. Your application code need not manage App Server sessions.

1.3 MarkLogic-Specific Key and Value Types

A MapReduce job configuration specifies the input and output key and value types for map and reduce. The MarkLogic Connector for Hadoop provides MarkLogic specific key and value classes, as follows:

| Class | Description |
|----------------------------|---|
| <code>DocumentURI</code> | A document URI. Use as a key type with <code>DocumentInputFormat</code> , <code>DocumentOutputFormat</code> , and <code>PropertyOutputFormat</code> . |
| <code>MarkLogicNode</code> | An XML node. Use as a value type with <code>DocumentInputFormat</code> , <code>NodeInputFormat</code> , <code>NodeOutputFormat</code> , and <code>PropertyOutputFormat</code> . |
| <code>NodePath</code> | A node URI. Use as a key type with <code>NodeInputFormat</code> and <code>NodeOutputFormat</code> . |

The MarkLogic Connector for Hadoop includes `InputFormat` and `OutputFormat` subclasses which predefine key-value pairs using the types listed above, such as `NodeInputFormat`. See “InputFormat Subclasses” on page 62 and “Using MarkLogic Server for Output” on page 66.

The MarkLogic Server specific types can be used in conjunction with non-connector types, such as `org.apache.hadoop.io.Text`. For example, a job using `NodeInputFormat` always has map input key-value pairs of type `(NodePath, MarkLogicNode)`, but can produce output key-value pairs of type `(Text, IntWritable)`.

For input data, you can also combine the MarkLogic Server specific types with certain Apache Hadoop MapReduce types in the same key-value pair by using `com.marklogic.mapreduce.KeyValueInputFormat` OR `com.marklogic.mapreduce.ValueInputFormat`. For details, see “Using `KeyValueInputFormat` and `ValueInputFormat`” on page 56.

The key and value types are usually configured programmatically through the `org.apache.hadoop.mapreduce.Job` API. For an example, see “InputFormat Subclasses” on page 62.

1.4 Deploying the Connector with a MarkLogic Server Cluster


This section covers the following topics:

- [Relationship of MarkLogic Server to a Hadoop Cluster](#)
- [Jobs Use In-Forest Evaluation](#)
- [Cluster-wide XDBC Configuration Requirements](#)

For more information about clustering, see [Clustering in MarkLogic Server](#) in the *Scalability, Availability, and Failover Guide*.

1.4.1 Relationship of MarkLogic Server to a Hadoop Cluster

Although it is possible to deploy Apache Hadoop MapReduce, HDFS, and MarkLogic Server on a single host for development purposes, production deployments usually involve a Hadoop cluster and a MarkLogic Server cluster, as shown below:



In a typical MapReduce/HDFS production deployment, a MapReduce Task Tracker runs on each data node host, though this is not required.

When using MarkLogic Server in a MapReduce job, whether or not to co-locate a Task Tracker with each MarkLogic Server node is dependent on your workload. Co-location reduces network traffic between the server and the MapReduce tasks, but places a heavier computational and memory burden on the host.

1.4.2 Jobs Use In-Forest Evaluation

To optimize performance, the MarkLogic Connector for Hadoop interacts with MarkLogic Server at the forest level. For example, if the reduce step of a job inserts a document into a MarkLogic Server database, the insert is an in-forest insert.

When MarkLogic Server is deployed across a cluster, the forests of a database can be distributed across multiple nodes. In this case, the in-forest evaluation of MapReduce related queries is also distributed across the MarkLogic Server nodes hosting the forests of the target database.

Therefore, every MarkLogic Server host that has at least one forest in a database used by a MapReduce job must be configured to act as both an e-node and a d-node. That is, each host must be capable of providing both query evaluation and data services. A pure d-node (for example, a host with a very small configured expanded tree cache) is not usable in a MapReduce job.

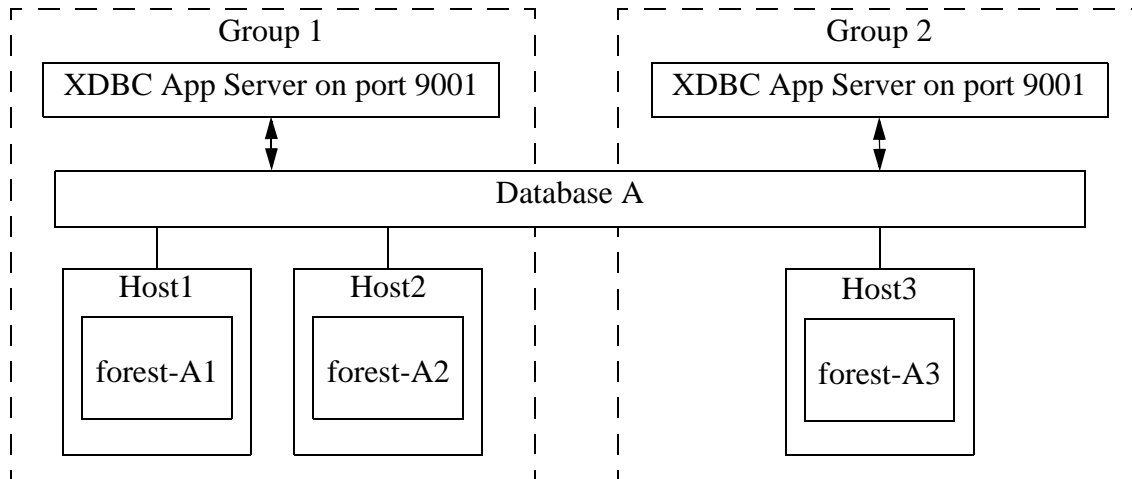
1.4.3 Cluster-wide XDBC Configuration Requirements

Because the MarkLogic Connector for Hadoop uses an XDBC App Server and in-forest query evaluation, your cluster might need special configuration to support MapReduce jobs.

If MarkLogic Server is used for input, each host that has at least one forest attached to the input database must have an XDBC App Server configured for that database. Additionally, the XDBC App Server must listen on the same port on each host. The same requirement applies to using MarkLogic Server for output. The input App Server, database and port can be the same or different from the output App Server, database and port.

Hosts within a group share the same App Server configuration, so you only need additional App Servers if hosts with forests attached to the input or output database are in multiple groups.

For example, the cluster shown below is properly configured to use Database A as a MapReduce input source. Database A has 3 forests, located on 3 hosts in 2 different groups. Therefore, both Group 1 and Group 2 must make Database A accessible on port 9001.



For details about the query evaluation, see “Jobs Use In-Forest Evaluation” on page 12. For information on related MapReduce job configuration properties, see “Identifying the Input MarkLogic Server Instance” on page 36 and “Identifying the Output MarkLogic Server Instance” on page 66.

1.5 Making a Secure Connection to MarkLogic Server with SSL

The MarkLogic Connector for Hadoop supports making secure connections to the input and output MarkLogic Server instances. To configure secure connections:

1. Enable SSL in the App Server, as described in [General Procedure for Setting up SSL for an App Server](#) in the *Administrator’s Guide*.
2. Create an implementation of the `com.marklogic.mapreduce.SslConfigOptions` interface in your job. Use one of the techniques described in [Accessing SSL-Enabled XDBC App Servers](#) in the *XCC Developer’s Guide* to provide a `javax.net.ssl.SSLContext` to the MarkLogic Connector for Hadoop.
3. Specify the `SslConfigOptions` subclass name from the previous step in the configuration property(s) `mapreduce.marklogic.input.ssloptionsclass` or `mapreduce.marklogic.output.ssloptionsclass`. See the examples below.
4. Enable SSL use by setting the configuration property(s) `mapreduce.marklogic.input.usessl` or `mapreduce.marklogic.output.usessl` to `true`.

You can set `mapreduce.marklogic.input.ssloptionsclass` and `mapreduce.marklogic.output.ssloptionsclass` either in a configuration file or programmatically. To set the property in a configuration file, set the value to your `SslConfigOptions` class name with “.class” appended to it. For example:

```
<property>
  <name>mapreduce.marklogic.input.ssloptionsclass</name>
```

```
<value>my.package.MySslOptions.class</value>
</property>
```

To set the property programatically, use the `org.apache.hadoop.conf.Configuration` API. For example:

```
import org.apache.hadoop.conf.Configuration;
import com.marklogic.mapreduce.SslConfigOptions;

public class ContentReader {
    static class MySslOptions implements SslConfigOptions {...}

    public static void main(String[] args) throws Exception {
        Configuration conf = new Configuration();
        ...

        conf.setClass("mapreduce.marklogic.input.ssloptionsclass",
            MySslOptions.class, SslConfigOptions.class);
        ...
    }
}
```

For a complete example of using SSL with the MarkLogic Connector for Hadoop, see “ContentReader” on page 103. For a basic XCC application using SSL, see [HelloSecureWorld](#) in the *XCC Developer’s Guide*.

2.0 Getting Started with the MarkLogic Connector for Hadoop

This chapter provides procedures for installing and configuring Apache Hadoop MapReduce and the MarkLogic Connector for Hadoop, and for running a simple MapReduce job that interacts with MarkLogic Server. For more examples, see “Using the Sample Applications” on page 87.

This chapter includes the following sections:

- [Requirements](#)
- [Installing Apache Hadoop and the MarkLogic Connector for Hadoop](#)
- [Configuring Hadoop to Use the Connector](#)
- [Running the HelloWorld Sample Application](#)
- [Making the Connector Available Across a Hadoop Cluster](#)
- [Organization of the Connector Distribution](#)

2.1 Requirements

This section covers the following topics:

- [Required Software](#)
- [Security Requirements for MapReduce Jobs](#)

2.1.1 Required Software

The MarkLogic Connector for Hadoop is a Java-only API and is only available on Linux. The MarkLogic Connector for Hadoop requires the following software:

- Linux
- MarkLogic MarkLogic Connector for Hadoop
- MarkLogic 5 or later
- MarkLogic XML Content Connector for Java (XCC/J) 5.0 or later
- Apache Hadoop MapReduce, version 0.20.2 only
- Java JDK 1.6 or later. The Oracle/Sun JDK is recommended.

Note: Apache Hadoop only supports the Oracle/Sun JDK, though other JDK’s may work. For details, see <http://wiki.apache.org/hadoop/HadoopJavaVersions>.

2.1.2 Security Requirements for MapReduce Jobs

The user with which a MapReduce job accesses MarkLogic Server must have appropriate privileges for the content accessed by the job, such as permission to read or update documents in the target database. Specify the user in the `mapreduce.marklogic.input.username` and `mapreduce.marklogic.output.username` job configuration properties. See “Configuring a MapReduce Job” on page 32.

In addition, the input and output user must use one of the pre-defined roles listed below:

| Role | Description |
|--------------------------------|--|
| <code>hadoop-user-read</code> | Enables use of MarkLogic Server as an input source for a MapReduce job. This role does not grant any other privileges, so the <code>mapreduce.marklogic.input.user</code> might still require additional privileges to read content from the target database. |
| <code>hadoop-user-write</code> | Enables use of MarkLogic Server as an output destination for a MapReduce job. This role does not grant any other privileges, so the <code>mapreduce.marklogic.output.user</code> might still require additional privileges to insert or update content in the target database. |
| <code>hadoop-user-all</code> | Combines the privileges of <code>hadoop-user-read</code> and <code>hadoop-user-write</code> . |

The `hadoop-internal` role is for internal use only. Do not assign this role to any users. This role is used to amp special privileges within the context of certain functions of the Hadoop MapReduce Connector. Assigning this role to users gives them privileges on the system that you typically do not want them to have.

For details about roles and privileges, see the *Understanding and Using Security Guide*.

2.2 Installing Apache Hadoop and the MarkLogic Connector for Hadoop

This section covers the following topics:

- [Installing Apache Hadoop](#)
- [Installing the MarkLogic Connector for Hadoop](#)

2.2.1 Installing Apache Hadoop

Skip this section if you already have an Apache Hadoop installation.

These instructions walk through configuring a single Apache Hadoop node, running in standalone mode. Standalone mode is the simplest configuration for development, but is not suitable for production. For other configurations, see “Deploying the Connector with a MarkLogic Server Cluster” on page 11 and the Apache Hadoop MapReduce installation instructions at:

http://hadoop.apache.org/common/docs/current/single_node_setup.html.

To install Apache Hadoop and configure it for standalone operation:

1. Download a supported version of Apache Hadoop MapReduce from hadoop.apache.org:

<http://hadoop.apache.org/mapreduce/releases.html>

For information on supported versions, see “Required Software” on page 15.

2. Unpack the Apache Hadoop zip file to a location of your choice. For example, if you download the Apache Hadoop zip file to `/space`, then unpack the file as follows:

```
$ cd /space
$ tar xzf hadoop-0.20.2.tar.gz
$ ls
hadoop-0.20.2/
hadoop-0.20.2.tar.gz
```

3. Optionally, set `HADOOP_HOME` in your shell environment to facilitate using the example commands in this guide. `$HADOOP_HOME` refers to the Hadoop installation location in the remainder of this guide. For example:

```
$ export HADOOP_HOME=/space/hadoop-0.20.2
```

4. If `JAVA_HOME` is not set in your environment by default, edit `$HADOOP_HOME/conf/hadoop-env.sh` and set `JAVA_HOME` to the root of your JDK installation. For example, if the JDK is installed as `/space/java/jdk1.6.0_26`, add the following near the top of `hadoop-env.sh`:

```
export JAVA_HOME=/space/java/jdk1.6.0_26
```

5. Run `hadoop` with no arguments to test your installation. You should see the Apache Hadoop usage message. For example:

```
$ $HADOOP_HOME/bin/hadoop
Usage: hadoop [--config confdir] COMMAND
...
```

Apache Hadoop is now usable in standalone mode.

2.2.2 Installing the MarkLogic Connector for Hadoop

This section assumes you are installing the Hadoop MapReduce Connector with the standalone Apache Hadoop installation created in “Installing Apache Hadoop” on page 16. For information about installation in an Apache Hadoop Cluster, see “Making the Connector Available Across a Hadoop Cluster” on page 23.

To install the MarkLogic Connector for Hadoop:

1. Download the MarkLogic Connector for Hadoop from developer.marklogic.com.
2. Unpack the connector package to a location of your choice. For example, assuming `/space/marklogic` contains the connector zip file and you install the MarkLogic Connector for Hadoop in `/space/marklogic/mapreduce`:

```
$ cd /space/marklogic
$ mkdir mapreduce; cd mapreduce
$ unzip ../MarkLogic-Connector-for-Hadoop-1.0.zip
```

3. Optionally, set `CONNECTOR_HOME` in your shell environment to facilitate using the example commands in this guide. The MarkLogic Connector for Hadoop installation directory is referred to as `$CONNECTOR_HOME` in this guide. For example:

```
$ export CONNECTOR_HOME=/space/marklogic/mapreduce
```

4. If XCC is not already installed, download XCC for Java from developer.marklogic.com and unzip the package to a location of your choice. The installation location, such as `/space/marklogic/xcc`, is referred to as `$XCC_HOME` in this guide.

```
$ cd /space/marklogic
$ mkdir xcc; cd xcc
$ unzip ../MarkXCC.Java-5.0.zip
```

Apache Hadoop must be configured to find the MarkLogic Connector for Hadoop libraries before you can use MarkLogic Server in a MapReduce job. See “Configuring Hadoop to Use the Connector” on page 18.

2.3 Configuring Hadoop to Use the Connector

Before using the MarkLogic Connector for Hadoop with your Apache Hadoop installation for the first time, place the MarkLogic Connector for Hadoop library and its dependent libraries on the Hadoop class path, as described below:

1. Open `$HADOOP_HOME/conf/hadoop-env.sh` for editing. Locate the `HADOOP_CLASSPATH` definition, if it already exists.
2. Modify the `HADOOP_CLASSPATH` definition, or add a definition, to add the MarkLogic Connector for Hadoop and XCC libraries to the `HADOOP_CLASSPATH`. Separate each path with a colon. For example, if using MarkLogic 5, the required libraries are:

- `$CONNECTOR_HOME/lib/commons-modeler-2.0.1.jar`
- `$CONNECTOR_HOME/lib/marklogic-mapreduce-1.0.jar`
- `$XCC_HOME/lib/marklogic-xcc-5.0.jar`

For example, if `$CONNECTOR_HOME` is `/space/marklogic/mapreduce` and `$XCC_HOME` is `/space/marklogic/xcc`, then `HADOOP_CLASSPATH` should be as follows (all on one line, with no whitespace):

```
export HADOOP_CLASSPATH=  
/space/marklogic/mapreduce/lib/commons-modeler-2.0.1.jar:  
/space/marklogic/mapreduce/lib/marklogic-mapreduce-1.0.jar:  
/space/marklogic/xcc/lib/marklogic-xcc-5.0.jar
```

Note: The version numbers in the jar file names can vary across MarkLogic Server releases.

Apache Hadoop MapReduce and the MarkLogic Connector for Hadoop are now ready for use.

2.4 Running the HelloWorld Sample Application

The section walks through configuring and running a simple HelloWorld sample job, assuming MarkLogic Server and Apache Hadoop are installed on the same single node, as described in “Installing Apache Hadoop and the MarkLogic Connector for Hadoop” on page 16.

The following steps are covered:

- [Creating the App Server and Database](#)
- [Loading the Input Data](#)
- [Configuring the Job](#)
- [Running the Job](#)

2.4.1 Creating the App Server and Database

Before running the sample application, create a database and an XDBC App Server. This guide assumes a database named “hello-hadoop” and an App Server on port 9001. For detailed instructions on creating databases and App Servers, see the *Administrator’s Guide*.

To create the database and App Server:

1. Open the Admin Interface to MarkLogic Server in your browser.
2. Create a forest named “hello-hadoop-1” using the default settings.
3. Create a database named “hello-hadoop” using the default settings.
4. Attach the “hello-hadoop-1” forest to the “hello-hadoop” database.

5. Go to the database configuration page for the “hello-hadoop” database and set the “directory creation” option to “manual”. The MarkLogic Connector for Hadoop requires this when inserting documents into the database.
6. Create an XDBC App Server named “hello-hadoop-xdbc”, attached to the hello-hadoop database. You may use any root and port number. This guide assumes these settings:

| | |
|-------------------------|--|
| xdbc server name | hello-hadoop-xdbc The XDBC server name. Required. You must supply |
| root | / The module directory root. Required. You must supply |
| port | 9001 The server socket bind internet Required. You must supply |
| modules | (file system) ▾ The database that contains app |
| database | hello-hadoop ▾ The database name. |

2.4.2 Loading the Input Data

Follow these instructions to initialize the “hello-hadoop” database with the input documents using Query Console. For details about Query Console, see the *Query Console User Guide*.

To load the database with the sample data:

1. Using your browser, launch Query Console on the MarkLogic Server instance to be used as an input source. For example, if the input XDBC App Server is running on `myhost`, visit this URL in the browser:

```
http://myhost:8000/qconsole
```

2. Create a new query in Query Console and replace the default contents with the following:

```
xquery version "1.0-ml";

let $hello := <data><child>hello mom</child></data>
let $world := <data><child>world event</child></data>

return(
  xdmp:document-insert("hello.xml", $hello),
  xdmp:document-insert("world.xml", $world)
)
```

3. In the `Content Source` dropdown, select the input XDBC App Server you configured for input in “Creating the App Server and Database” on page 19.
4. Select `Text` as the output format and click `Run` to execute the query.
5. Click the `Explore` button at the top of the query editor to examine the database contents. You should see `hello.xml` and `world.xml` in the database.

2.4.3 Configuring the Job

Before running the `HelloWorld` sample job, set the connector configuration properties that identify the MarkLogic Server user and instance for input and output. Although the input and output MarkLogic Server instances and users can be different, this example configures the job to use the host and port of the `hello-hadoop-xdbc` App Server created in “Creating the App Server and Database” on page 19 for both input and output.

Configuration also includes an input and an output user name. Choose or create a user with sufficient privileges to access the `hello-hadoop-xdbc` App Server, and read and insert documents in the `hello-hadoop` database. If using a non-admin user, assign the user to the `hadoop-user-all` role. For details, see “Security Requirements for MapReduce Jobs” on page 16.

To configure the job:

1. Copy the `marklogic-hello-world.xml` configuration file from `$CONNECTOR_HOME/conf` to `$HADOOP_HOME/conf`. For example:

```
$ cp $CONNECTOR_HOME/conf/marklogic-hello-world.xml $HADOOP_HOME/conf
```

2. Edit `$HADOOP_HOME/conf/marklogic-hello-world.xml` to configure your input and output host name, port, user name, and password. Set the following parameters to match your environment:

```
mapreduce.marklogic.input.username
mapreduce.marklogic.input.password
mapreduce.marklogic.input.host
mapreduce.marklogic.input.port
mapreduce.marklogic.output.username
mapreduce.marklogic.output.password
mapreduce.marklogic.output.host
mapreduce.marklogic.output.port
```

The configured input user must have sufficient privileges to access the XDBC App Server identified by the input host/port and to read documents from the input database.

The configured output user must have sufficient privileges to access the XDBC App Server identified by the output host/port and to insert documents in the output database.

2.4.4 Running the Job

The `HelloWorld` sample reads the first word of text from the input documents, concatenates the words into a string, and saves the result as `HelloWorld.txt`. Assuming the database contains only the documents created in “Loading the Input Data” on page 20, the output document contains the phrase “hello world”. If your database contains additional documents, you get different results.

To view the sample code, see `$CONNECTOR_HOME/src/com/marklogic/mapreduce/examples`.

Use the following command line to run the `HelloWorld` job. Modify the version number in the connector JAR file name, as needed.

```
$HADOOP_HOME/bin/hadoop jar \  
  $CONNECTOR_HOME/lib/marklogic-mapreduce-examples-1.0.jar \  
  com.marklogic.mapreduce.examples>HelloWorld
```

As the job runs, Hadoop reports the job progress to `stdout`. If the sample job does not run or does not produce the expected results, see “Troubleshooting and Debugging” on page 85.

Near the end of the job output, you should see text similar to the following. Notice there are 2 map input records (`hello.xml` and `world.xml`), 2 map output records (the first word from each input record), and 1 reduce output record (`HelloWorld.txt`).

```
INFO mapred.JobClient: map 100% reduce 100%  
INFO mapred.JobClient: Job complete: job_local_0001  
INFO mapred.JobClient: Counters: 12  
INFO mapred.JobClient:   FileSystemCounters  
INFO mapred.JobClient:     FILE_BYTES_READ=56115  
INFO mapred.JobClient:     FILE_BYTES_WRITTEN=101467  
INFO mapred.JobClient:   Map-Reduce Framework  
INFO mapred.JobClient:     Combine output records=0  
INFO mapred.JobClient:     Map input records=2  
INFO mapred.JobClient:     Reduce shuffle bytes=0  
INFO mapred.JobClient:     Reduce output records=1  
INFO mapred.JobClient:     Spilled Records=4  
INFO mapred.JobClient:     Map output bytes=20  
INFO mapred.JobClient:     Combine input records=0  
INFO mapred.JobClient:     Map output records=2  
INFO mapred.JobClient:     Reduce input records=2
```

Use Query Console to explore the output database and examine the output document, `HelloWorld.txt`. The document should contain the phrase “hello world”.

If you do not see the expected output, see the tips in “Troubleshooting and Debugging” on page 85.

2.5 Making the Connector Available Across a Hadoop Cluster

When you submit a MapReduce job to run on an Apache Hadoop cluster, the job resources must be accessible by the master Job Tracker node and all worker nodes. Job resources include the job JAR file, configuration files, and all dependent libraries. When you use the MarkLogic Connector for Hadoop in your job, this includes the connector and XCC JAR files.

You must always have the job resources available on the Hadoop node where you launch the job. Dependent JAR files, such as the MarkLogic Connector for Hadoop libraries must be on the `HADOOP_CLASSPATH` on this node, as described in “Configuring Hadoop to Use the Connector” on page 18.

Hadoop offers many options for making job resources available to the worker nodes, including:

- Using the `-libjars` Hadoop command line option and parsing the options in your main class using `org.apache.hadoop.util.GenericOptionsParser`.
- Bundling dependent libraries and other resources into your job JAR file.
- Storing dependent libraries and other resources in HDFS or other shared file system and using the Apache Hadoop DistributedCache to locate and load them.
- Installing required software on all nodes in the cluster.

The best solution depends upon the needs of your application and environment. See the Apache Hadoop documentation for more details on making resources available across a Hadoop cluster.

2.6 Organization of the Connector Distribution

The MarkLogic Connector for Hadoop distribution has the following layout:

| Document or Directory | Description |
|-----------------------|--|
| <code>conf/</code> | The XML config files for the sample applications. For details, see “Using the Sample Applications” on page 87. |

| Document or Directory | Description |
|-----------------------|--|
| docs/ | The Javadoc for the connector in both expanded HTML and compressed zip format. |
| lib/ | The connector and connector examples jar files, <code>marklogic-mapreduce-1.0.jar</code> and <code>marklogic-mapreduce-examples-1.0.jar</code> . Note that the jar file names include the version number, so the names in your installation might be slightly different. |
| src/ | The source code for the sample applications. |
| Readme.txt | Includes the connector version number and any last-minute updates not included in the documentation. |

3.0 Apache Hadoop MapReduce Concepts

This chapter provides a very brief introduction to Apache Hadoop MapReduce. If you are already familiar with Apache Hadoop MapReduce, skip this chapter. For a complete discussion of the MapReduce and the Hadoop framework, see the Hadoop documentation, available from the Apache Software Foundation:

<http://hadoop.apache.org/common/docs/current/>

This chapter covers the following topics:

- [MapReduce Overview](#)
- [Example: Calculating Word Occurrences](#)
- [Understanding the MapReduce Job Life Cycle](#)
- [How Hadoop Partitions Map Input Data](#)
- [Configuring a MapReduce Job](#)
- [Running a MapReduce Job](#)
- [Viewing Job Status and Logs](#)

3.1 MapReduce Overview

Apache Hadoop MapReduce is a framework for processing large data sets in parallel across a Hadoop cluster. Data analysis uses a two step map and reduce process. The job configuration supplies map and reduce analysis functions and the Hadoop framework provides the scheduling, distribution, and parallelization services.


The top level unit of work in MapReduce is a job. A job usually has a map and a reduce phase, though the reduce phase can be omitted. For example, consider a MapReduce job that counts the number of times each word is used across a set of documents. The map phase counts the words in each document, then the reduce phase aggregates the per-document data into word counts spanning the entire collection.

During the map phase, the input data is divided into input splits for analysis by map tasks running in parallel across the Hadoop cluster. By default, the MapReduce framework gets input data from the Hadoop Distributed File System (HDFS). Using the MarkLogic Connector for Hadoop enables the framework to get input data from a MarkLogic Server instance. For details, see “Map Task” on page 30.

The reduce phase uses results from map tasks as input to a set of parallel reduce tasks. The reduce tasks consolidate the data into final results. By default, the MapReduce framework stores results in HDFS. Using the MarkLogic Connector for Hadoop enables the framework to store results in a MarkLogic Server instance. For details, see “Reduce Task” on page 30.

Although the reduce phase depends on output from the map phase, map and reduce processing is not necessarily sequential. That is, reduce tasks can begin as soon as any map task completes. It is not necessary for all map tasks to complete before any reduce task can begin.

MapReduce operates on key-value pairs. Conceptually, a MapReduce job takes a set of input key-value pairs and produces a set of output key-value pairs by passing the data through map and reduce functions. The map tasks produce an intermediate set of key-value pairs that the reduce tasks uses as input. The diagram below illustrates the progression from input key-value pairs to output key-value pairs at a high level:



Though each set of key-value pairs is homogeneous, the key-value pairs in each step need not have the same type. For example, the key-value pairs in the input set (KV1) can be `(string, string)` pairs, with the map phase producing `(string, integer)` pairs as intermediate results (KV2), and the reduce phase producing `(integer, string)` pairs for the final results (KV3). See “Example: Calculating Word Occurrences” on page 27.

The keys in the map output pairs need not be unique. Between the map processing and the reduce processing, a shuffle step sorts all map output values with the same key into a single reduce input `(key, value-list)` pair, where the “value” is a list of all values sharing the same key. Thus, the input to a reduce task is actually a set of `(key, value-list)` pairs.

The key and value types at each stage determine the interfaces to your map and reduce functions. Therefore, before coding a job, determine the data types needed at each stage in the map-reduce process. For example:

1. Choose the reduce output key and value types that best represents the desired outcome.
2. Choose the map input key and value types best suited to represent the input data from which to derive the final result.
3. Determine the transformation necessary to get from the map input to the reduce output, and choose the intermediate map output/reduce input key value type to match.

Control MapReduce job characteristics through configuration properties. The job configuration specifies:

- how to gather input
- the types of the input and output key-value pairs for each stage
- the map and reduce functions

- how and where to store the final results

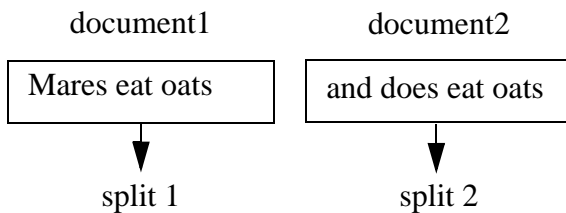
For more information on job configuration, see “Configuring a MapReduce Job” on page 32. For information on MarkLogic specific configuration properties, see “Input Configuration Properties” on page 60 and “Output Configuration Properties” on page 79.

3.2 Example: Calculating Word Occurrences

This example demonstrates the basic MapReduce concept by calculating the number of occurrence of each each word in a set of text files. For an in-depth discussion and source code for an equivalent example, see the Hadoop MapReduce tutorial at:


http://hadoop.apache.org/mapreduce/docs/current/mapred_tutorial.html

Recall that MapReduce input data is divided into input splits, and the splits are further divided into input key-value pairs. In this example, the input data set is the two documents, `document1` and `document2`. The `InputFormat` subclass divides the data set into one split per document, for a total of 2 splits:



A `(line number, text)` key-value pair is generated for each line in an input document. The map function discards the line number and produces a per-line `(word, count)` pair for each word in the input line. The reduce phase produces `(word, count)` pairs representing aggregated word counts across all the input documents.

Given the input data shown above, the map-reduce progression for the example job is:



The output from the map phase contains multiple key-value pairs with the same key: The “oats” and “eat” keys appear twice. Recall that the MapReduce framework consolidates all values with the same key before entering the reduce phase, so the input to reduce is actually `(key, values)` pairs. Therefore, the full progression from map output, through reduce, to final results is:

3. Copies the job resources (configuration, job JAR file, input splits) to a shared location, such as an HDFS directory, where it is accessible to the Job Tracker and Task Trackers.
4. Submits the job to the Job Tracker.

3.3.2 Job Tracker

The Job Tracker is responsible for scheduling jobs, dividing a job into map and reduce tasks, distributing map and reduce tasks among worker nodes, task failure recovery, and tracking the job status. Job scheduling and failure recovery are not discussed here; see the Apache Hadoop MapReduce documentation.

When preparing to run a job, the Job Tracker:

1. Fetches input splits from the shared location where the Job Client placed the information.
2. Creates a map task for each split.
3. Assigns each map task to a Task Tracker (worker node).

The Job Tracker monitors the health of the Task Trackers and the progress of the job. As map tasks complete and results become available, the Job Tracker:

1. Creates reduce tasks up to the maximum enabled by the job configuration.
2. Assigns each map result partition to a reduce task.
3. Assigns each reduce task to a Task Tracker.

A job is complete when all map and reduce tasks successfully complete, or, if there is no reduce step, when all map tasks successfully complete.

3.3.3 Task Tracker

A Task Tracker manages the tasks of one worker node and reports status to the Job Tracker. Often, the Task Tracker runs on the associated worker node, but it is not required to be on the same host.

When the Job Tracker assigns a map or reduce task to a Task Tracker, the Task Tracker:

1. Fetches job resources locally.
2. Spawns a child JVM on the worker node to execute the map or reduce task.
3. Reports status to the Job Tracker.

The task spawned by the Task Tracker runs the job's map or reduce functions.

3.3.4 Map Task

The Hadoop MapReduce framework creates a map task to process each input split. The map task:

1. Uses the `InputFormat` to fetch the input data locally and create input key-value pairs.
2. Applies the job-supplied map function to each key-value pair.
3. Performs local sorting and aggregation of the results.
4. If the job includes a `Combiner`, runs the `Combiner` for further aggregation.
5. Stores the results locally, in memory and on the local file system.
6. Communicates progress and status to the Task Tracker.

Map task results undergo a local sort by key to prepare the data for consumption by reduce tasks. If a `Combiner` is configured for the job, it also runs in the map task. A `Combiner` consolidates the data in an application-specific way, reducing the amount of data that must be transferred to reduce tasks. For example, a `Combiner` might compute a local maximum value for a key and discard the rest of the values. The details of how map tasks manage, sort, and shuffle results are not covered here. See the Apache Hadoop MapReduce documentation.

When a map task notifies the Task Tracker of completion, the Task Tracker notifies the Job Tracker. The Job Tracker then makes the results available to reduce tasks.

3.3.5 Reduce Task

The reduce phase aggregates the results from the map phase into final results. Usually, the final result set is smaller than the input set, but this is application dependent. The reduction is carried out by parallel reduce tasks. The reduce input keys and values need not have the same type as the output keys and values.

Note: The reduce phase is optional. You may configure a job to stop after the map phase completes. For details, see “Configuring a Map-Only Job” on page 34.

Reduce is carried out in three phases, copy, sort, and merge. A reduce task:

1. Fetches job resources locally.
2. Enters the copy phase to fetch local copies of all the assigned map results from the map worker nodes.
3. When the copy phase completes, executes the sort phase to merge the copied results into a single sorted set of `(key, value-list)` pairs.

4. When the sort phase completes, executes the reduce phase, invoking the job-supplied reduce function on each `(key, value-list)` pair.
5. Saves the final results to the output destination, such as HDFS.

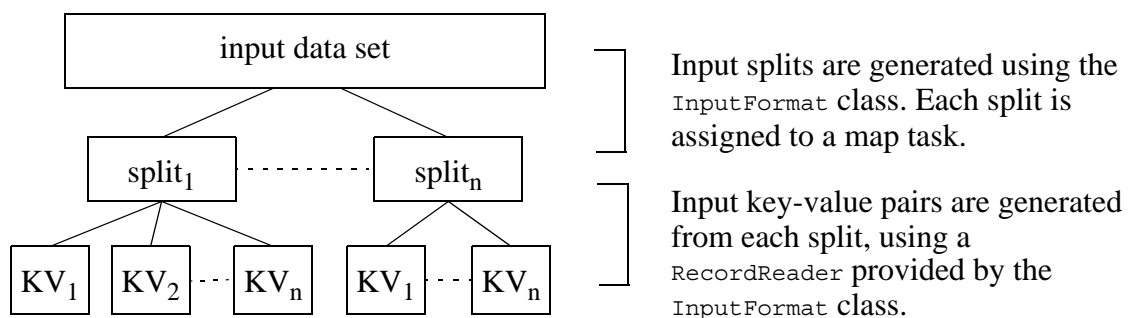
The input to a reduce function is key-value pairs where the value is a list of values sharing the same key. For example, if one map task produces a key-value pair `(“eat”, 2)` and another map task produces the pair `(“eat”, 1)`, then these pairs are consolidated into `(“eat”, (2, 1))` for input to the reduce function. If the purpose of the reduce phase is to compute a sum of all the values for each key, then the final output key-value pair for this input is `(“eat”, 3)`. For a more complete example, see “Example: Calculating Word Occurrences” on page 27.

Output from the reduce phase is saved to the destination configured for the job, such as HDFS or MarkLogic Server. Reduce tasks use an `OutputFormat` subclass to record results. Apache Hadoop provides `OutputFormat` subclasses for using HDFS as the output destination. The MarkLogic Connector for Hadoop provides `OutputFormat` subclasses for using a MarkLogic Server database as the destination. For a list of available subclasses, see “OutputFormat Subclasses” on page 83. The connector also provides classes for defining key and value types; see “MarkLogic-Specific Key and Value Types” on page 10.

3.4 How Hadoop Partitions Map Input Data

When you submit a job, the MapReduce framework divides the input data set into chunks called *splits* using the `org.apache.hadoop.mapreduce.InputFormat` subclass supplied in the job configuration. Splits are created by the local Job Client and included in the job information made available to the Job Tracker.

The JobTracker creates a map task for each split. Each map task uses a `RecordReader` provided by the `InputFormat` subclass to transform the split into input key-value pairs. The diagram below shows how the input data is broken down for analysis during the map phase:



Apache Hadoop provides `InputFormat` subclasses for using HDFS as an input source. The MarkLogic Connector for Hadoop provides `InputFormat` subclasses for using MarkLogic Server as an input source. For a list of available MarkLogic-specific subclasses, see “InputFormat Subclasses” on page 62.

3.5 Configuring a MapReduce Job

This section covers the following topics:

- [Configuration Basics](#)
- [Setting Properties in a Configuration File](#)
- [Setting Properties Using the Hadoop API](#)
- [Setting Properties on the Command Line](#)
- [Configuring a Map-Only Job](#)

3.5.1 Configuration Basics

Apache Hadoop configuration is controlled by multiple layers of configuration files and property settings. You may set configuration properties in configuration files, programmatically, and on the command line. For details, see the Apache Hadoop MapReduce documentation at:

http://hadoop.apache.org/common/docs/current/cluster_setup.html#Configuration+Files

Configuration properties can include the following:

- Hadoop MapReduce properties
- Application-specific properties, such as the properties defined by the MarkLogic Connector for Hadoop API in `com.marklogic.mapreduce.MarkLogicConstants`
- Job-specific properties. See the `mapreduce.linkcount.baseuri` property used by the `LinkCountInDoc` sample application

A MapReduce application must configure at least the following:

- **Mapper:** Define a subclass of `org.apache.hadoop.mapreduce.Mapper`, usually overriding at least the `Map.map()` method.
- **InputFormat:** Select a subclass of `org.apache.hadoop.mapreduce.InputFormat` and pass it to `org.apache.hadoop.mapreduce.Job.setInputFormatClass`.

If the job includes a reduce step, then the application must also configure the following:

- **Reducer:** Define a subclass of `org.apache.hadoop.mapreduce.Reducer`, usually overriding at least the `Reducer.reduce()` method.
- **OutputFormat:** Select a subclass of `org.apache.hadoop.mapreduce.OutputFormat` and pass it to `org.apache.hadoop.mapreduce.Job.setOutputFormatClass`.
- **Output key and value types for the map and reduce phases.** Set the key and value types appropriate for your `InputFormat` and `OutputFormat` subclasses using the Job API functions, such as `org.apache.hadoop.mapreduce.Job.setMapOutputKeyClass`.

For details about configuring MapReduce jobs, see the Apache Hadoop MapReduce documentation. For details about connector-specific configuration options, see “Using MarkLogic Server for Input” on page 36 and “Using MarkLogic Server for Output” on page 66.

3.5.2 Setting Properties in a Configuration File

Configuration files are best suited for static configuration properties. By default, Apache Hadoop looks for configuration files in `$HADOOP_HOME/conf`. You may override this location on the `hadoop` command line. For details, see the Apache Hadoop Commands Guide:

http://hadoop.apache.org/common/docs/r0.20.2/commands_manual.html

Job configuration files are XML files with the following layout:

```
<?xml version="1.0"?>
<?xml-stylesheet type="text/xsl" href="configuration.xsl?>
<configuration>
  <property>
    <name>the.property.name</name>
    <value>the.property.value</value>
  </property>
  <property>...</property>
</configuration>
```

3.5.3 Setting Properties Using the Hadoop API

You can use the Apache Hadoop API to set properties that cannot be set in a configuration file or that have dynamically calculated values. For example, you can set the `InputFormat` class by calling `org.apache.hadoop.mapreduce.Job.setInputFormatClass`.

Set properties prior to submitting the job. That is, prior to calling

`org.apache.mapreduce.Job.submit` OR `org.apache.mapreduce.Job.waitForCompletion`.

To set an arbitrary property programmatically, use the `org.apache.hadoop.conf.Configuration` API. This API includes methods for setting property values to string, boolean, numeric, and class types. Set the properties prior to starting the job. For example, to set the MarkLogic Connector for Hadoop `marklogic.mapreduce.input.documentselector` property, at runtime, you can do the following:

```
import org.apache.hadoop.conf.Configuration;
import org.apache.hadoop.mapreduce.Job;
import com.marklogic.mapreduce.MarkLogicConstants;
...
public class myJob { ...
  public static void main(String[] args) throws Exception {
    Configuration conf = new Configuration();
    String[] otherArgs =
      new GenericOptionsParser(conf, args).getRemainingArgs();
    Job job = new Job(conf);
```

```
// Build up the document selector dynamically...
String mySelector = ...;

conf = job.getConfiguration();
conf.set(MarkLogicConstants.DOCUMENT_SELECTOR, mySelector);
...
System.exit(job.waitForCompletion(true) ? 0 : 1);
}
}
```

3.5.4 Setting Properties on the Command Line

You may set properties on the hadoop command line using the `-D` command line option. For example:

```
$ hadoop -Dproperty_name=value jar myJob.jar myClass
```

3.5.5 Configuring a Map-Only Job

Some jobs can complete all their work during the map phase. For example, the `ContentLoader` sample application, which loads documents from HDFS into MarkLogic Server, performs document loading in the map tasks and then stops without using reduce tasks. See “`ContentLoader`” on page 104.

To stop a job after the map completes, set the number of reduce tasks to zero by setting the property `mapred.reduce.tasks` to 0 in your job configuration file, or by calling `org.apache.hadoop.Configuration.setNumReduceTasks(0)` in your application.

3.6 Running a MapReduce Job

Use the `hadoop jar` command to execute a job. For example:

```
$ hadoop jar /path/myJob.jar my.package.Class options
```

For specific examples, see “Using the Sample Applications” on page 87. For details about the hadoop command line, see the Apache Hadoop MapReduce Command Guide at:

http://hadoop.apache.org/common/docs/current/commands_manual.html

3.7 Viewing Job Status and Logs

A running job reports progress and errors to `stdout` and `stderr`. If Apache Hadoop is configured for standalone mode, this output is all that is available. If Apache Hadoop is configured in pseudo-distributed or fully distributed mode, logs are recorded in `$HADOOP_HOME/logs` by default. You can also view job status and results using the Job Tracker and NameNode web consoles.

Use the Job Tracker web console to view job status. By default, the Job Tracker console is available on port 50030. For example:

```
http://localhost:50030
```

Use the NameNode web console to browse HDFS, including job results, and to look at job related logs. By default, the NameNode console is available on port 50070. For example:

```
http://localhost:50070
```

To examine logs for a job:

1. Navigate to the NameNode console on port 50070.
2. Click the Namenode Logs link at the top of the page. A listing of the logs directory appears.
3. To check for errors, click on a Job Tracker log file, such as `hadoop-your_username-jobtracker-your_hostname.log`. The contents of the log file appears.
4. Click on your browser Back button to return to the log directory listing.
5. To view output output from your job, click on userlogs at the bottom of the logs directory listing. A user log listing appears.
6. Locate the run corresponding to your job by looking at the timestamps. A log directory is created for each map task and each reduce task.

Map tasks have `_m_` in the log directory name. Reduce tasks have `_r_` in the log directory name. For example, `attempt_201111020908_0001_m_000001_0/` represents a map task.
7. Click on the map or reduce task whose output you want to examine. A directory listing of the stdout, stderr, and syslog for the select task.
8. Click on the name of a log to see what the task output to that destination. For example, click stderr to view text written to stderr by the task.

For details, see the Apache Hadoop MapReduce documentation.

4.0 Using MarkLogic Server for Input

When using MarkLogic Server as an input source, the input key-value pairs passed to the map function of a map task are constructed by transforming database fragments or lexicon data in the input split into key-value pairs using the configured `InputFormat` subclass and input properties. This section covers the following topics:

- [Basic Steps](#)
- [Basic Input Mode](#)
- [Advanced Input Mode](#)
- [Using KeyValueInputFormat and ValueInputFormat](#)
- [Configuring a Map-Only Job](#)
- [Input Configuration Properties](#)
- [InputFormat Subclasses](#)

4.1 Basic Steps

To configure the map phase of a MapReduce job to use MarkLogic Server for input, perform the following steps:

- [Identifying the Input MarkLogic Server Instance](#)
- [Specifying the Input Mode](#)
- [Specifying the Input Key and Value Types](#)
- [Defining the Map Function](#)

4.1.1 Identifying the Input MarkLogic Server Instance

The MarkLogic Server input instance is identified by setting job configuration properties. For general information on setting configuration properties, see “Configuring a MapReduce Job” on page 32.

Specify the following properties to identify the input MarkLogic Server instance:

| Property | Description |
|---|---|
| <code>mapreduce.marklogic.input.host</code> | Hostname or IP address of the server hosting your input XDBC App Server. The host must be resolvable by the nodes in your Hadoop cluster, so you should usually not use “local-host”. |
| <code>mapreduce.marklogic.input.port</code> | The port configured for the target XDBC App Server on the input host. |
| <code>mapreduce.marklogic.input.username</code> | Username privileged to read from the database attached to your XDBC App Server. |
| <code>mapreduce.marklogic.input.password</code> | Cleartext password for the <code>input.username</code> user. |

When you use MarkLogic Server in a cluster, all MarkLogic Server hosts containing a forest in the input database must be accessible through an XDBC server on the same port. The host identified in the configuration properties may be any qualifying host in the cluster. For details, see “Deploying the Connector with a MarkLogic Server Cluster” on page 11.

You can configure a job to connect to the App Server through SSL by setting the `mapreduce.marklogic.input.usessl` property. For details, see “Making a Secure Connection to MarkLogic Server with SSL” on page 13. For an example, see “ContentReader” on page 103.

For more information on the properties, see “Input Configuration Properties” on page 60.

4.1.2 Specifying the Input Mode

The MarkLogic Connector for Hadoop input mode determines how much responsibility your job has for creating input splits and input key-value pairs. The MarkLogic Connector for Hadoop supports basic and advanced input modes. Basic mode is the default. Set the input mode using the `mapreduce.marklogic.input.mode` configuration property.

When MarkLogic Server is the input source, each map task runs an input query against the task input split to select the fragments or records from which to create map input key-value pairs. An *input split query* divides the input content into input splits by forest, and an *input query* selects content within the split. For a general discussion of input splits in MapReduce, see “How Hadoop Partitions Map Input Data” on page 31.

In basic mode, the MarkLogic Connector for Hadoop uses a built-in input split query and builds the input query based on data selection properties defined by your job. This enables the connector to optimize the interaction between the Apache Hadoop MapReduce framework and MarkLogic Server. If your input selection needs can be met by the basic mode query construction properties, you should use basic mode as it offers the best performance. For details, see “Basic Input Mode” on page 40.

Basic mode supports selecting input data from documents or a lexicon. You can only use one of these methods in a job. If configuration properties are set for both methods, the connector uses a lexicon for input. If none of the input selection properties are set, the connector uses the default document selectors.

In advanced mode, your application provides the input split query and input query. Using advanced mode gives you complete control over map input key-value pair creation, but adds complexity. For details, see “Advanced Input Mode” on page 50.

4.1.3 Specifying the Input Key and Value Types

As discussed in “How Hadoop Partitions Map Input Data” on page 31, the `org.apache.hadoop.mapreduce.InputFormat` subclass configured for the job determines the types of the input keys and values and how the content selected by the input query is transformed into key-value pairs.

To specify the `InputFormat` subclass programmatically, use the `org.apache.hadoop.mapreduce.Job` API. The following example configures the job to use `NodeInputFormat`, which creates `(NodePath, MarkLogicNode)` input key-value pairs:

```
import com.marklogic.mapreduce.NodeInputFormat;
import org.apache.hadoop.mapreduce.Job;
...
public class LinkCountInDoc {
    ...
    public static void main(String[] args) throws Exception {
        Job job = new Job(conf);
        job.setInputFormatClass(NodeInputFormat.class);
        ...
    }
}
```

The MarkLogic Connector for Hadoop API includes several `InputFormat` subclasses for MarkLogic Server data types, such as `(document URI, node)` key-value pairs. For details, see “InputFormat Subclasses” on page 62.

You can also use Apache Hadoop MapReduce key-value types through the `KeyValueInputFormat` and `ValueInputFormat` subclasses. These classes define type conversions between MarkLogic Server types and standard Apache Hadoop MapReduce types; see “Using `KeyValueInputFormat` and `ValueInputFormat`” on page 56.

4.1.4 Defining the Map Function

A MapReduce application must include a subclass of `org.apache.hadoop.mapreduce.Mapper` and implement a `Mapper.map()` method. The `map` method transforms the map input key-value pairs into output key-value pairs that can be used as input for the reduce step.

The application-specific `Mapper` subclass and the signature of the `Mapper.map` method must match the configured map phase input and output key-value pair types. For example, if the map phase uses `NodeInputFormat` for input and produces `(Text, IntWritable)` output key-value pairs, then the `Mapper` subclass should be similar to the following because `NodeInputFormat` creates `(NodePath, MarkLogicNode)` input key-value pairs:

```
import org.apache.hadoop.mapreduce.Mapper;
import org.apache.hadoop.io.Text;
import org.apache.hadoop.io.IntWritable;
import com.marklogic.mapreduce.MarkLogicNode;
import com.marklogic.mapreduce.NodePath;

public class LinkCountInDoc {
    ...
    public static class RefMapper
        extends Mapper<NodePath, MarkLogicNode, Text, IntWritable> {
        public void map(NodePath key, MarkLogicNode value, Context context)
        {
            ...derive output key(s) and value(s)...
            context.write(output_key, output_value)
        }
    }
}
```

The `Mapper.map()` method constructs result key-value pairs corresponding to the expected output type, and then writes pairs to the `org.apache.hadoop.mapreduce.Context` parameter for subsequent handling by the MapReduce framework. In the example above, the map output key-value pairs must be `(Text, IntWritable)` pairs.

Configure the `Mapper` into the job using the `org.apache.hadoop.mapreduce.Job` API. For example:

```
public static void main(String[] args) throws Exception {
    Configuration conf = new Configuration();
    ...
    Job job = new Job(conf);
    job.setMapperClass(RefMapper.class);
    ...
}
```

For a list of `InputFormat` subclasses provided by the MarkLogic Connector for Hadoop API and the associated key and value types, see “[InputFormat Subclasses](#)” on page 62.

For more details, see “Example: Counting Href Links” on page 48 and the sample code provided in the connector package.

4.2 Basic Input Mode

The MarkLogic Connector for Hadoop supports basic and advanced input modes through the `mapreduce.marklogic.input.mode` configuration property. The default mode, basic, offers the best performance. In basic input mode, the connector handles all aspects of split creation. The job configuration properties control which fragments in a split to transform into input key-value pairs.

This section covers the following topics:

- [Creating Input Splits](#)
- [Using a Lexicon to Generate Key-Value Pairs](#)
- [Using XPath to Generate Key-Value Pairs](#)
- [Example: Counting Href Links](#)

For details on advanced input mode, see “Advanced Input Mode” on page 50.


4.2.1 Creating Input Splits

Basic mode does not give you control over split creation, other than setting the maximum split size using the `mapreduce.marklogic.input.maxsplitsize` property. Basic mode does give the job control over the content passed to the map function, as described in “Specifying the Input Mode” on page 37.

When using document fragments for input data, the connector divides the fragments in the database into input splits by forest, such that each split contains at most the number of fragments specified by `mapreduce.marklogic.input.maxsplitsize`.

Multiple splits might be required to cover a single forest. A split never spans more than one forest. When the maximum split size is smaller than the total number of fragments in the forest, the MarkLogic Connector for Hadoop may adjust the split size downwards so that the size is balanced across a given forest.

For example, consider a database containing 2 forests, forest1 and forest2. Forest1 contains 15 documents. Forest2 contains 18 documents. If `input.maxsplitsize` is 10, then the connector creates 2 splits from each forest, with each split covering roughly half the documents in each forest:



A similar distribution of data occurs when using a lexicon for input. Splits are constructed by querying the lexicon in each forest and assigning at most `input.maxsplitsize` entries to each split.

Use advanced mode if you require control over split creation. See “Advanced Input Mode” on page 50.

4.2.2 Using a Lexicon to Generate Key-Value Pairs

This section describes how to use a MarkLogic Server lexicon to create map input key-value pairs. Using a lexicon precludes using an XPath document selector for input. For general information about lexicons, see [Browsing With Lexicons](#) in the *Search Developer’s Guide*.

To use a lexicon function:

1. [Implement a Lexicon Function Wrapper Subclass](#) corresponding to the XQuery lexicon function you wish to use.
2. [Override Lexicon Function Parameter Wrapper Methods](#) to specify the parameter values for the lexicon function call.
3. [Choose an InputFormat](#) subclass.
4. [Configure the Job](#) to use the lexicon by setting the configuration property `mapreduce.marklogic.input.lexiconfunctionclass`.
5. Additional [De-duplication of Results Might Be Required](#) if your application requires uniqueness in the map output values.

4.2.2.1 Implement a Lexicon Function Wrapper Subclass

The MarkLogic Connector for Hadoop API package `com.marklogic.mapreduce.functions` includes a lexicon function abstract class corresponding to each supported XQuery API lexicon function. Range and geospatial lexicons are not supported at this time. For a mapping between lexicon wrapper classes and the XQuery lexicon functions, see “Lexicon Function Subclasses” on page 64.

The subclass you create encapsulates a call to an XQuery lexicon function. The lexicon function is implicit in the parent class, and the function call parameters are implemented by the methods. For example, to use `cts:element-attribute-value-co-occurrences`, implement a subclass of `com.marklogic.mapreduce.functions.ElemAttrValueCooccurrences`:

```
import com.marklogic.mapreduce.functions.ElemAttrValueCooccurrences;

public class LinkCountCooccurrences {
    static class HrefTitleMap extends ElemAttrValueCooccurrences {...}
}
```

In your subclass, override the inherited methods for the required parameters and for the optional parameters to be included in the call, as described in “Override Lexicon Function Parameter Wrapper Methods” on page 42.

4.2.2.2 Override Lexicon Function Parameter Wrapper Methods

Each lexicon function wrapper class includes methods for defining the supported function parameter values, as strings. The connector uses this information to construct a call to the wrapped lexicon function in the input query.

The parameter related methods vary by the function interface. For details on a particular wrapper class, see the MarkLogic Connector for Hadoop javadoc.

You may not specify values corresponding to the `$quality-weight` or `$forest-ids` parameters of any lexicon function. Quality weight is not used in a MapReduce context, and the connector manages forest constraints for you in basic mode.

Note: If you include options settings in the lexicon call, do not set the `skip` or `truncate` options. The MarkLogic Connector for Hadoop reserves the `skip` and `truncate` lexicon function options for internal use.

Note: The `frequency-order` option supported by some lexicon calls is not honored in MapReduce results.

For example, the XQuery lexicon function `cts:element-values` is exposed through the `com.marklogic.mapreduce.function.ElementValues` class. The prototype for `cts:element-values` is:

```

cts:element-values (
  $element-names as xs:QName*,
  [$start as xs:anyAtomicType?],
  [$options as xs:string*],
  [$query as cts:query?],
  [$quality-weight as xs:double?],
  [$forest-ids as xs:unsignedLong*]
)

```

The `com.marklogic.mapreduce.function.ElementValues` wrapper class includes abstract methods corresponding to the `$element-names`, `$start`, `$options`, and `$query` parameters:

```

package com.marklogic.mapreduce.functions;

public abstract class LexiconFunction {

    public String getLexiconQuery() {...}
    public String[] getUserDefinedOptions() {...}
}

public abstract class ValuesOrWordsFunction extends LexiconFunction {
    public String getStart() {...}
}

public abstract class ElementValues extends ValuesOrWordsFunction {
    public abstract String[] getElementNames();
}

```

The parameter method overrides must return the string representation of an XQuery expression that evaluates to the expected parameter type. For example, if your job needs to make a lexicon function call equivalent to the following, then it must override the methods related to the `$element-names` and `$options` parameters to `cts:element-values`:

```

cts:element-values(xs:QName("wp:a"), (), "ascending")

```

The default start value from the inherited `ValuesOrWordsFunction.getStart` method already matches the desired parameter value. Override `ElementValue.getElementNames` and `LexiconFunction.getUserDefinedOptions` to specify the other parameter values. The resulting subclass is:

```

import com.marklogic.mapreduce.functions.ElementValues;
...
public class myLexiconFunction extends ElementValues {
    public String[] getElementNames() {
        String[] elementNames = {"xs:QName(\"wp:a\")"};
        return elementNames;
    }

    public String[] getUserDefinedOptions() {
        String[] options = {"ascending"};
        return options;
    }
}

```

```
    }
  }
```

Note: If you override `ValuesOrWordsFunction.getStart` to specify a start value of type `xs:string`, you must include escaped double-quotes in the returned string. For example, to specify a start value of "aardvark", the `getStart` override must return `"\"aardvark\""`.

For a complete example, see `com.marklogic.mapreduce.examples.LinkCountCooccurrences`, included with the MarkLogic Connector for Hadoop.

4.2.2.3 Choose an InputFormat

Lexicon functions return data from a lexicon, rather than document content, so the document oriented `InputFormat` classes such as `DocumentInputFormat` and `NodeInputFormat` are not applicable with lexicon input. Instead, use `com.marklogic.mapreduce.ValueInputFormat` or `com.marklogic.mapreduce.KeyValueInputFormat`.

Note: `KeyValueInputFormat` is only usable with co-occurrences lexicon functions. `ValueInputFormat` may be used with any lexicon function.

When using `ValueInputFormat` with a lexicon function, choose an input value type that either matches the type returned by the lexicon function or can be converted from the lexicon return type to the input value type. For details on `KeyValueFormat` and `ValueInputFormat` and the supported type conversions, see “Using `KeyValueInputFormat` and `ValueInputFormat`” on page 56.

For example, `cts:element-values` returns `xs:anyAtomicType*`. If the element range index queried by your `ElementValues` subclass is an index with scalar type `int`, then you might configure the job to use `IntWritable` as the map input value type.

The co-occurrences lexicon function classes generate key-value pairs corresponding to the `cts:value` elements in each `cts:co-occurrence` returned by the underlying lexicon function. The key is the first `cts:value` in each `cts:co-occurrence`, and the value is the second. For example, if `cts:element-values` returns the following XML:

```
<cts:co-occurrence xmlns:cts="http://marklogic.com/cts"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <cts:value xsi:type="xs:string">MARCELLUS</cts:value>
  <cts:value xsi:type="xs:string">BERNARDO</cts:value>
</cts:co-occurrence>
<cts:co-occurrence xmlns:cts="http://marklogic.com/cts"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <cts:value xsi:type="xs:string">ROSENCRANTZ</cts:value>
  <cts:value xsi:type="xs:string">GUILDENSTERN</cts:value>
</cts:co-occurrence>
<cts:co-occurrence xmlns:cts="http://marklogic.com/cts"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
```

```

    xmlns:xs="http://www.w3.org/2001/XMLSchema">
    <cts:value xsi:type="xs:string">HORATIO</cts:value>
    <cts:value xsi:type="xs:string">MARCELLUS</cts:value>
  </cts:co-occurrence>

```

Then the generated map input key-value pairs are:

```

("MARCELLUS", "BERNARDO")
("ROSENCRANTZ", "GUILDENSTERN")
("HORATIO", "MARCELLUS")

```

As with `valueInputFormat`, choose key and value types corresponding to the lexicon type or convertible from the lexicon type. For a complete example, see `com.marklogic.mapreduce.examples.LinkCountCooccurrences`.

4.2.2.4 Configure the Job

Set the `mapreduce.marklogic.input.lexiconfunctionclass` job configuration property to specify which lexicon call to use. This property setting takes precedence over `mapreduce.marklogic.input.documentselector` and `mapreduce.marklogic.input.subdocumentexpr`.

You can set `mapreduce.marklogic.input.lexiconfunctionclass` either in a configuration file or programmatically. To set the property in a configuration file, set the value to your lexicon subclass name with “.class” appended to it. For example:

```

<property>
  <name>mapreduce.marklogic.input.lexiconfunctionclass</name>
  <value>my.package.LexiconFunction.class</value>
</property>

```

To set the property programatically, use the `org.apache.hadoop.conf.Configuration` API. For example:

```

import org.apache.hadoop.conf.Configuration;
import com.marklogic.mapreduce.functions.ElemAttrValueCooccurrences;

public class LinkCountCooccurrences {
    static class HrefTitleMap extends ElemAttrValueCooccurrences {...}

    public static void main(String[] args) throws Exception {
        Configuration conf = new Configuration();
        ...

        conf.setClass(MarkLogicConstants.INPUT_LEXICON_FUNCTION_CLASS,
            HrefTitleMap.class, ElemAttrValueCooccurrences.class);
        ...
    }
}

```

4.2.2.5 De-duplication of Results Might Be Required

Your job might need to de-duplicate map results if the analysis performed by your job depends upon unique values in the map output key-value pairs. Only a URI lexicon is guaranteed to have unique values across all forests.

MarkLogic Server maintains lexicon data per forest. Each forest includes a lexicon for the fragments in that forest. Normally, when you directly call an XQuery lexicon function, MarkLogic Server gathers results from each forest level lexicon and de-duplicates them. As discussed in “Creating Input Splits” on page 40, the MarkLogic Connector for Hadoop runs input queries directly against each forest. The in-forest evaluation circumvents the de-duplication of results from lexicon functions.

For example, consider a database with 2 forests, each containing one document. If the database is configured with an element word lexicon for the `<data>` element, then the table below reflects the contents of the word lexicon in each forest:

| forest | document content | words in per-forest lexicon |
|----------|--|-----------------------------|
| sample-1 | <code><data></code> hello world <code></data></code> | hello world |
| sample-2 | <code><data></code> goodbye world <code></data></code> | goodbye world |

Calling `cts:words` (the lexicon function underlying `com.marklogic.mapreduce.functions.Words`) against the database returns (“goodbye”, “hello”, “world”). Therefore, if you use the corresponding lexicon function class in a MapReduce job, you might expect 3 map output key-value pairs from this input data.

However, the in-forest evaluation of `cts:words` used by the MarkLogic Connector for Hadoop results in the two word lists (“hello”, “world”) and (“goodbye”, “world”), generating 4 map output key-value pairs. If the purpose of the job is to count the number of unique words in the content, simply counting the number of reduce input key-value pairs results in an incorrect answer of 4.

If you require this kind of uniqueness, you must de-duplicate the results in your application. For example, your Reducer might use a hash map to detect and discard duplicate values.

4.2.3 Using XPath to Generate Key-Value Pairs

This section describes how to use XPath components to generate map input key-value pairs from document fragments. Using an XPath document selector precludes using a lexicon for input.

Use the following configuration properties to select the input fragments in a split based on an XPath expression:

- `mapreduce.marklogic.input.documentselector`
- `mapreduce.marklogic.input.subdocumentexpr`

The MarkLogic Connector for Hadoop constructs an input query that includes a path expression equivalent to concatenating the document selector and sub-document expression. For example, if the document selector is `fn:collection()` and the sub-document expression is `/*:data`, then the constructed input query uses an input XPath expression equivalent to `fn:collection/*:data`.

The document selector determines which documents within a forest are included in each split. The sub-document expression determines which fragments of the selected documents contribute input key-value pairs. The MarkLogic Connector for Hadoop API uses separate document selector and sub-document expression components to maximize internal input query optimization.

Note: The document selector must be a partially searchable XPath expression. For details, see [Fully Searchable Paths and cts:search Operations](#) in the *Query Performance and Tuning Guide*.

The default document selector is `fn:collection()`, which selects all documents in the split. The default sub-document expression returns all nodes selected by the document selector. If you do not specify either property, all documents nodes in the split are used.

The table below shows the results from several example document selector and subdocument expression combinations.

| document selector | subdocument expression | result |
|---|---------------------------|--|
| none | none | All document nodes in the split. |
| <code>fn:collection("wikipedia")</code> | none | All document nodes in the split that are in the “wikipedia” collection. |
| <code>fn:collection("wikipedia")</code> | <code>//wp:nominee</code> | All <code>wp:nominee</code> elements in documents in the split that are in the “wikipedia” collection. |

If your document selector or subdocument expression uses namespace qualifiers, define the namespaces using the `mapreduce.marklogic.input.namespaces` configuration property. The property value is a comma-separated list of alias-URI pairs. The following example defines a “wp” namespace and then uses it in the subdocument expression:

```

<property>
  <name>mapreduce.marklogic.input.namespace</name>
  <value>wp, http://www.marklogic.com/wikipedia</value>
</property>
<property>
  <name>mapreduce.marklogic.input.subdocumentexpr</name>
  <value>wp:nominee</value>
</property>

```

For details, see “Example: Counting Href Links” on page 48 and the sample code included in the connector package.

4.2.4 Example: Counting Href Links

Consider the `LinkCountInDoc` example, described in “LinkCountInDoc” on page 98. This example counts hyperlinks within the input collection. The map function for this sample expects (node URI, node) input pairs and produces (href title, href count) output pairs.

Suppose the database contains an input document in the “wikipedia” collection, with the URI `/oscars/drama.xml` and the following content (ellided for brevity):

```

<wp:nominee>...
  <wp:film>...
    <wp:p>...
      <wp:a href="http://en.wikipedia.org/wiki/Drama_film"
        title="Drama film">
        This is a dramatic film
      </wp:a>
    </wp:p>
  </wp:film>
</wp:nominee>

```

Notice the content includes embedded hrefs with titles. The following job configuration property settings extract hrefs with titles that reference other Wikipedia articles:

```

<property>
  <name>mapreduce.marklogic.input.namespace</name>
  <value>wp, http://www.marklogic.com/wikipedia</value>
</property>
<property>
  <name>mapreduce.marklogic.input.documentselector</name>
  <value>fn:collection("wikipedia")</value>
</property>
<property>
  <name>mapreduce.marklogic.input.subdocumentexpr</name>
  <value>//wp:a[@href and @title and fn:starts-with(@href,
    "http://en.wikipedia.org")]@title</value>
</property>

```

Using the above document selector and subdocument expression, the connector constructs an input query that first selects documents in the `split` that are in the “wikipedia” collection. Then, from those documents, all title attribute nodes in `wp:a` elements that contain hrefs to Wikipedia articles. Here is a breakdown of the path expression properties that drive the input query;

```
documentselector: fn:collection("wikipedia")
subdocumentexpr:
  //wp:a[                               (: all anchors :)
  @href and                             (: with href attributes :)
  @title and                             (: and title attributes :)
  fn:starts-with(@href, "http://en.wikipedia.org")]
                                          (: referring to Wikipedia articles :)
  /@title                               (: return the title attribute node :)
```

The configured `InputFormat` subclass controls transformation of the fragments selected by the input query into input key-value pairs. For example, the input query for the sample document shown above finds a matching node with these characteristics:

```
document URI: /oscars/drama.xml
node path: /wp:nominee/wp:film/wp:abstract/html:p[1]/html:a[1]/@title
node: attribute node for title="Drama film" from this anchor element:
<wp:a href="http://en.wikipedia.org/wiki/Drama_film"
      title="Drama film">
  drama film
</wp:a>
```

If you configure the job to use `NodeInputFormat`, the input key-value pairs are (`NodePath`, `MarkLogicNode`) pairs. That is, the key is the node URI and the value is the node. The sample data shown above results in this input pair:

```
key: /wp:nominee/wp:film/wp:p[1]/wp:a[1]/@title
value: attribute node for title="Drama film"
```

If you configure the job to use `DocumentInputFormat` instead, the input key-value pairs have type (`DocumentURI`, `MarkLogic`). The sample data results in the following input pair, differing from the previous case only in the key:

```
key: /oscars/drama.xml
value: attribute node for title="Drama film"
```

For a list of `InputFormat` subclasses provided by the connector, include the types of keys and values generated by each, see “`InputFormat` Subclasses” on page 62.

4.3 Advanced Input Mode

The connector supports basic and advanced input modes through the `mapreduce.marklogic.input.mode` configuration property. Advanced mode gives your application complete control over input split creation and fragment selection, at the cost of added complexity. In advanced mode, you must supply an input split query and an input query which the connector uses to map phase input.

This section covers the following topics:

- [Creating Input Splits](#)
- [Creating Input Key-Value Pairs](#)
- [Optimizing Your Input Query](#)
- [Example: Counting Hrefs Using Advanced Mode](#)

For details on basic mode, see “Basic Input Mode” on page 40. For a general discussion of input splits and key-value pair creation, see “How Hadoop Partitions Map Input Data” on page 31.

4.3.1 Creating Input Splits

This section covers the following topics:

- [Overview](#)
- [Creating a Split Query with `hadoop:get-splits`](#)

4.3.1.1 Overview

In advanced mode, your application controls input split creation by supplying an input split query in the `mapreduce.marklogic.input.splitquery` configuration property. Your input split query must generate `(forest-id, record-count, host-name)` tuples. Build a split query using the XQuery library function `hadoop:get-splits`, in combination with your own XPath and `cts:query`.

The split query returns a host name and forest id because the MarkLogic Connector for Hadoop interacts with MarkLogic Server at the forest level. When a split is assigned to a map task, the connector running in the task submits the input query directly against the forest identified in the split tuple. The `host-name` and `forest-id` in the split tuple identify the target forest.

The `record-count` in the split tuple only needs to be an rough estimate of the number of input key-value pairs in the split. The estimate need not be accurate. What constitutes a record is job-dependent. For example, a record can be a document fragment, a node, or a value from a lexicon.

For example, basic input mode uses a simple estimate of the total number of documents in each forest. When the input query runs against a split, it can generate more or fewer input key-value pairs than estimated. The more accurate the record count estimate is, the more accurately Hadoop balances workload across tasks.

An input split never spans multiple forests, but the content in a single forest may span multiple splits. The maximum number of records in a split is determined by the `com.marklogic.mapreduce.maxsplitsize` configuration property. The connector, rather than your split query, handles bounding splits by this maximum.

For example, if the split query returns a count of 1000 fragments for a forest and the max split size is 600, the connector generates two splits for that forest, one for the first 500 fragments and the other for the next 500 fragments. The connector adjusts the actual split size downward as needed to generate splits of similar size, so you do not get one split of 600 fragments and another of 400 fragments. This rebalancing keeps the workload even across tasks.

4.3.1.2 Creating a Split Query with `hadoop:get-splits`

Use the XQuery library function `hadoop:get-splits` to create split tuples using a searchable expression and `cts:query`. The parameters to `hadoop:get-splits` determine the documents under consideration in each forest, equivalent to the `$expression` and `$query` parameters of `cts:search`. The function returns an estimate rather than a true count to improve performance.

The following split query example returns one tuple per forest, where the count is an estimate of the total number of documents in the forest.

```
<property>
  <name>mapreduce.marklogic.input.splitquery</name>
  <value><![CDATA[
    declare namespace wp="http://www.mediawiki.org/xml/export-0.4/";
    import module namespace hadoop =
      "http://marklogic.com/xdmp/hadoop" at "/MarkLogic/hadoop.xqy";
    hadoop:get-splits('', 'fn:doc()', 'cts:and-query(())')
  ]]></value>
</property>
```

You may write an input split query without `hadoop:get-splits`, but your job may require additional privileges if you do so. The example below is equivalent to the previous `hadoop:get-splits` example. Unlike the previous example, however, admin privileges are required to execute this sample query.

```
<property>
  <name>mapreduce.marklogic.input.splitquery</name>
  <value><![CDATA[
    declare namespace wp="http://marklogic.com/wikipedia";
    import module namespace admin = "http://marklogic.com/xdmp/admin"
      at "/MarkLogic/admin.xqy";
    let $conf := admin:get-configuration()
    for $forest in xdmp:database-forests(xdmp:database())
      let $host_id := admin:forest-get-host($conf, $forest)
      let $host_name := admin:host-get-name($conf, $host_id)
      let $cnt :=
        xdmp:estimate(
          cts:search(fn:doc(), cts:and-query(()), (), 0.0, $forest))
    return ($forest, $cnt, $host_name)
  ]]></value>
</property>
```

```

    ]]></value>
  </property>

```

4.3.2 Creating Input Key-Value Pairs

As described in “Understanding the MapReduce Job Life Cycle” on page 28, map input key-value pairs are generated by each map task from the input split assigned to the task. The MarkLogic Server content in the split is selected by the input query in the `mapreduce.marklogic.input.query` configuration property.

The input query must return a sequence, but the nature of the items in the sequence is dependent on the configured `InputFormat` subclass. For example, if the job uses `NodeInputFormat`, the input key-value pairs are `(NodePath, MarkLogicNode)` pairs. Therefore, an input query for `NodeInputFormat` must return a sequence of nodes. The MarkLogic Connector for Hadoop then converts each node into a `(NodePath, MarkLogicNode)` pair for passage to the map function.

The table below summarizes the input query results expected by each `InputFormat` subclass provided by the MarkLogic Connector for Hadoop. For more information about the classes, see “InputFormat Subclasses” on page 62.

| InputFormat subclass | Input query result |
|----------------------------------|--|
| <code>DocumentInputFormat</code> | <code>document-node()*</code> |
| <code>NodeInputFormat</code> | <code>node()*</code> |
| <code>KeyValueInputFormat</code> | A sequence of alternating keys and values, <code>(key1, value1, ..., keyN, valueN)</code> . The key and value types depend on the job configuration. See “Using <code>KeyValueInputFormat</code> and <code>ValueInputFormat</code> ” on page 56. |
| <code>ValueInputFormat</code> | A sequence of values, <code>(value1, ..., valueN)</code> . The value type depends on the job configuration. See “Using <code>KeyValueInputFormat</code> and <code>ValueInputFormat</code> ” on page 56. |

`KeyValueInputFormat` and `ValueInputFormat` enable you to configure arbitrary key and/or value types, within certain constraints. For details, see “Using `KeyValueInputFormat` and `ValueInputFormat`” on page 56. When using these types, your input query must return keys and/or values compatible with the type conversions defined in “Supported Type Transformations” on page 58.

The following example uses `ValueInputFormat` with `org.apache.hadoop.mapreduce.Text` as the value type. (The key type with `ValueInputFormat` is always `org.apache.hadoop.mapreduce.LongWritable`, and is supplied automatically by the MarkLogic Connector for Hadoop.) The input query returns an attribute node, and the implicit type conversion built into `ValueInputFormat` converts the result into the attribute text for passage to the map function.

```
<property>
  <name>mapreduce.marklogic.input.query</name>
  <value><![CDATA[
    declare namespace wp="http://marklogic.com/wikipedia";
    for $t in fn:collection()/wp:nominee/wp:a[@title and @href]/@title
      return $t
  ]]></value>
</property>
```

4.3.3 Optimizing Your Input Query

In basic mode, the MarkLogic Connector for Hadoop optimizes the input query built from the configuration properties. In advanced input mode, you must optimize the input query yourself. Use the configuration property `mapreduce.marklogic.input.bindsplitrange` to optimize your input query by limiting the work your query does per split.

Usually, each input split covers a set of records (or fragments) in a specific forest; see the illustration in “Creating Input Splits” on page 40. The input split might cover only a subset of the content in a forest, but the input query runs against the entire forest. Constraining the content manipulated by the input query to just those under consideration for the split can significantly improve job performance.

For example, imagine an input query built around the following XPath expression:

```
fn:collection()//wp:a[@title]
```

In the worst case, evaluating the above expression examines every document in the forest for each split. If the forest contains 100 documents, covered by 4 splits of 25 fragments each, then the job might examine $4 * 100$ documents to cover the content in the forest, and then discard all but 25 of the results in each split:

```
fn:collection()//wp:a[@title] [1 to 25]
fn:collection()//wp:a[@title] [26 to 50]
fn:collection()//wp:a[@title] [51 to 75]
fn:collection()//wp:a[@title] [76 to 100]
```

The range on each expression above, such as `[1 to 25]`, is the split range. Constraining the set of documents by the split range earlier can be more efficient:

```
fn:collection() [1 to 25]//wp:a[@title]
...
```

In practice MarkLogic Server might internally optimize such a simple XPath expression, but a more complex FLOWR expression can require hand tuning.

The exact optimization is query dependent, but the split start and end values are made available to your input query through external variables if you set the configuration property `mapreduce.marklogic.input.bindsplitrange` to `true`.

To use this feature, set the property to true and do the following in your input query:

1. Declare the pre-defined namespace `http://marklogic.com/hadoop`. The split start and end variables are in this namespace. For example:

```
declare namespace mlmr="http://marklogic.com/hadoop";
```

2. Declare `splitstart` and `splitend` as external variables in the namespace declared in Step 1. You must use only these names. For example:

```
declare variable $mlmr:splitstart as xs:integer external;
declare variable $mlmr:splitend as xs:integer external;
```

3. Use the variables to constrain your query.

For example, the following input query returns the first 1K bytes of each document in a collection of binary documents. Notice that query uses `splitstart` and `splitend` to construct a sub-binary node for just the documents relevant to each split:

```
xquery version "1.0-ml";
declare namespace mlmr="http://marklogic.com/hadoop";           (: 1 :)

declare variable $mlmr:splitstart as xs:integer external;      (: 2 :)
declare variable $mlmr:splitend as xs:integer external;

for $doc in fn:doc()[$mlmr:splitstart to $mlmr:splitend]       (: 3 :)
return xdmp:subbinary($doc/binary(), 0, 1000)
```

For a complete example, see “BinaryReader” on page 102.

4.3.4 Example: Counting Hrefs Using Advanced Mode

This example counts hyperlinks with titles using advanced input mode. The full code is available in the `LinkCount` example included in the MarkLogic Connector for Hadoop package. See “LinkCount” on page 100.

Suppose the input data has the following structure:

```
<wp:page>
  <title>The Topic</title>
  ...
  <text>
```

```

    <p>...
    <a href="Drama film" title="Drama film">
      Drama film
    </a>
  </p>
</text>
</wp:page>

```

Notice the content includes embedded hrefs with titles. The input split query shown below estimates the number of documents in each forest by calling `xdmp:estimate` and returns (forest-id, count, host-name) tuples. The query is wrapped in a CDATA section so the contents are not interpreted by the Apache Hadoop MapReduce configuration parser.

```

<property>
  <name>mapreduce.marklogic.input.splitquery</name>
  <value><![CDATA[
    declare namespace wp="http://www.mediawiki.org/xml/export-0.4/";
    import module namespace admin =
      "http://marklogic.com/xdmp/admin" at "/MarkLogic/admin.xqy";
    let $conf := admin:get-configuration()
    for $forest in xdmp:database-forests(xdmp:database())
      let $host_id := admin:forest-get-host($conf, $forest)
      let $host_name := admin:host-get-name($conf, $host_id)
      let $cnt :=
        xdmp:estimate(
          cts:search(fn:doc(), cts:and-query(), (), 0.0, $forest))
      return ($forest, $cnt, $host_name)
  ]]></value>
</property>

```

The split query generates at least one split per forest. More than one split can be generated per forest, depending on the number of documents in the forest and the value of the configuration property `mapreduce.marklogic.input.maxsplitsize`. For details, see “Creating Input Splits” on page 50.

The example input query selects qualifying hrefs from each document in the split. A qualifying href in this example is one that has a title and refers to another Wikipedia topic, rather than to an image or an external file. The input query below uses an XPath expression to find qualifying hrefs and extract the text from the title attribute:

```

<property>
  <name>mapreduce.marklogic.input.query</name>
  <value><![CDATA[
    xquery version "1.0-ml";
    declare namespace wp="http://www.mediawiki.org/xml/export-0.4/";
    //wp:a[@title and @href and not
      (fn:starts-with(@href, "#")
      or fn:starts-with(@href, "http://")
      or fn:starts-with(@href, "File:")
      or fn:starts-with(@href, "Image:"))]/@title
  ]]></value>

```

```
    ]]></value>
  </property>
```

The `LinkCount` sample uses the above split query and input query with `ValueInputFormat` configured to generate `(LongWritable, Text)` map input pairs, where the key is a unique (but uninteresting) number and the value is the text from the title attribute of qualifying hrefs. For example one input to the map function might be:

```
(42, "Drama film")
```

For more information about `ValueInputFormat`, see “Using `KeyValueInputFormat` and `ValueInputFormat`” on page 56.

The map function of `LinkCount` discards the input key and generates a `(Text, IntWritable)` key-value pair where the key is the text from the title attribute (the value on the input pair), and the value is always one, to indicate a single reference to the topic:

```
public static class RefMapper
  extends Mapper<LongWritable, Text, Text, IntWritable> {
    private final static IntWritable one = new IntWritable(1);
    private Text refURI = new Text();

    public void map(LongWritable key, Text value, Context context)
      throws IOException, InterruptedException
    {
      refURI.set(value);
      context.write(refURI, one);
    }
    ...
}
```

For example, if `RefMapper.map` receives `(42, "Drama film")` as input, then it produces the output pair `("Drama film", 1)`.

4.4 Using `KeyValueInputFormat` and `ValueInputFormat`

The `KeyValueInputFormat` and `ValueInputFormat` classes enable you to create input key-value pairs with types other than the connector specific types `MarkLogicNode`, `NodePath`, and `DocumentURI`.

This section covers the following topics about using `KeyValueInputFormat` and `ValueInputFormat`.

- [Overview](#)
- [Job Configuration](#)
- [Supported Type Transformations](#)
- [Example: Using `KeyValueInputFormat`](#)

4.4.1 Overview

Use `KeyValueInputFormat` or `ValueInputFormat` to specify your own key-value type combinations, or when extracting input data from a MarkLogic Server lexicon. The connector performs a “best-effort” type conversion between the key and/or value types from the input query or lexicon function and the target type. For details, see “Supported Type Transformations” on page 58.

Use `KeyValueInputFormat` type when your application depends upon both the key and value type. Use `ValueInputFormat` when only the value is interesting. In `ValueInputFormat` the key is simply a unique number with no significance.

Most lexicon functions require use of `ValueInputFormat`. You may use `KeyValueInputFormat` only with lexicon co-occurrence functions. For details, see “Using a Lexicon to Generate Key-Value Pairs” on page 41.

For an example, see “Example: Using `KeyValueInputFormat`” on page 58.

4.4.2 Job Configuration

When using `KeyValueInputFormat`, specify the map phase input key and value type by setting the configuration properties `mapreduce.marklogic.input.keyclass` and `mapreduce.marklogic.input.valueclass`. For example, to set the properties in the configuration file:

```
<property>
  <name>mapreduce.marklogic.input.keyclass</name>
  <value>org.apache.hadoop.io.Text.class</value>
</property>
<property>
  <name>mapreduce.marklogic.input.valueclass</name>
  <value>com.marklogic.mapreduce.NodePath.class</value>
</property>
```

You can also use the `org.apache.hadoop.mapreduce.Job` API to set the property in code:

```
import org.apache.hadoop.io.Text;
import org.apache.hadoop.io.Writable;
...
public static void main(String[] args) throws Exception {
    Configuration conf = new Configuration();
    ...
    conf.setClass("mapreduce.marklogic.input.valueClass",
                 Text.class, Writable.class);
}
```

When using `ValueInputFormat`, specify the map phase value type by setting the configuration property `mapreduce.marklogic.input.valueclass`. The key type is built in and is always `org.apache.hadoop.io.LongWritable`.

4.4.3 Supported Type Transformations

The set of supported type transformations are listed in the table below. The list applies to both the key and value types when using `KeyValueInputFormat`, and to the value type when using `ValueInputFormat`.

| If the target key type or value type is | Then the query result type must be |
|--|---|
| <code>org.apache.hadoop.io.Text</code> | Any XQuery type with a <code>String</code> representation. See <code>com.marklog.xcc.types.XdmValue.asString()</code> . |
| <code>org.apache.hadoop.io.IntWritable</code> <code>org.apache.hadoop.io.VIntWritable</code> <code>org.apache.hadoop.io.LongWritable</code> <code>org.apache.hadoop.io.VLongWritable</code> | <code>xs:integer</code> |
| <code>org.apache.hadoop.io.BooleanWritable</code> | <code>xs:boolean</code> |
| <code>org.apache.hadoop.io.FloatWritable</code> | <code>xs:float</code> |
| <code>org.apache.hadoop.io.DoubleWritable</code> | <code>xs:double</code> |
| <code>org.apache.hadoop.io.BytesWritable</code> | <code>xs:hexBinary</code> <code>xs:base64Binary</code> |
| <code>com.marklogic.mapreduce.MarkLogicNode</code> | <code>node()</code> |

If the query result type and target type do not correspond to one of the supported conversions, or the target type is `Text` but the result type does not have a `String` representation, `java.lang.UnsupportedOperationException` is raised.

4.4.4 Example: Using KeyValueInputFormat

Consider a job which uses `KeyValueInputFormat` to count the number of occurrences of each href in the input document set and saves the results to HDFS. This is similar to the `LinkCount` sample included with the connector; see `com.marklogic.mapreduce.examples.LinkCount.java` and the `conf/marklogic-advanced.xml` configuration file.

This example uses advanced input mode with an input query that generates (`xs:string`, `xs:integer`) pairs containing the href title and a count from `cts:frequency`. The input query is shown below:

```
declare namespace wp="http://www.mediawiki.org/xml/export-0.4/";
declare variable $M := cts:element-attribute-value-co-occurrences (
  xs:QName ("wp:a"),
  xs:QName ("href"),
  xs:QName ("wp:a"),
  xs:QName ("title"),
```

```

("proximity=0", "map",
 "collation=http://marklogic.com/collation/codepoint"),
cts:directory-query("/space/wikipedia/enwiki/", "infinity")) ;
for $k in map:keys($M) [
  not(starts-with(., "#") or starts-with(., "http://") or
    starts-with(., "File:") or starts-with(., "Image:")) ]
let $v := map:get($M, $k)
where $v != ""
return
  for $each in $v
    return ($each, cts:frequency($each))

```

KeyValueInputFormat generates (Text, IntWritable) map phase input key-value pairs from the above query. The key is the title from the href anchor and the value is the count. The following input configuration properties define the key and value types for KeyValueInputFormat:

```

<property>
  <name>mapreduce.marklogic.input.keyclass</name>
  <value>org.apache.hadoop.io.Text</value>
</property>
<property>
  <name>mapreduce.marklogic.input.valueclass</name>
  <value>org.apache.hadoop.io.IntWritable</value>
</property>

```

The job programmatically sets the InputFormat class to KeyValueFormat, and the Mapper.map method expects the corresponding input types, as shown below. KeyValueInputFormat handles converting the input query results for each pair from (xs:string, xs:integer) into (Text, IntWritable).

```

import org.apache.hadoop.mapreduce.Job;
import org.apache.hadoop.mapreduce.Mapper;
import com.marklogic.mapreduce.KeyValueInputFormat;

public class LinkCount {
  public static class RefMapper
    extends Mapper<Text, IntWritable, Text, IntWritable> {
    public void map(Text key, IntWritable value, Context context)
    {...}
    ...

  public static void main(String[] args) throws Exception {
    Configuration conf = new Configuration();
    ...
    Job job = new Job(conf);
    job.setInputFormatClass(KeyValueInputFormat.class);
    ...
  }
}

```

The same type conversions apply when using `ValueInputFormat`, but you do not need to configure `mapreduce.marklogic.input.keyclass` because the key class is always `LongWritable`. The `LinkCount` sample code uses `ValueInputFormat` and relies on the default key and value types to achieve the same effect as shown above.

4.5 Configuring a Map-Only Job

Some jobs do not require a reduce step. For example, a job which ingests documents in HDFS into MarkLogic Server may complete the ingestion during the map step. For an example of a map-only job, see “ContentLoader” on page 104.

To stop a job after the map completes, set the number of reduce tasks to zero. You may set the number of reduce tasks in a configuration file or programmatically. For example, to set the number of reduce tasks to 0 in a configuration, include the following setting:

```
<property>
  <name>mapred.reduce.tasks</name>
  <value>0</value>
</property>
```

To set the number of reduce tasks to 0 programmatically, use the Hadoop API function `org.apache.hadoop.Configuration.setNumReduceTasks`.

If your map-only job uses MarkLogic Server for output, you must disable the Apache Hadoop MapReduce speculative execution feature. For details, see “Disabling Speculative Execution” on page 69.

4.6 Input Configuration Properties

The table below summarizes connector configuration properties for using MarkLogic Server as an input source. Some properties may only be used with basic input mode and others may only be used with advanced input mode. For more information, see `com.marklogic.mapreduce.MarkLogicConstants` in the *MarkLogic Connector for Hadoop API Reference*.

| Property | Description |
|---|---|
| <code>mapreduce.marklogic.input.username</code> | Username privileged to read from the database attached to your XDBC App Server. |
| <code>mapreduce.marklogic.input.password</code> | Cleartext password for the <code>input.username</code> user. |
| <code>mapreduce.marklogic.input.host</code> | Hostname of the server hosting your input XDBC App Server. |
| <code>mapreduce.marklogic.input.port</code> | The port configured for your XDBC App Server on the input host. |
| <code>mapreduce.marklogic.input.usessl</code> | Whether or not to use an SSL connection to the server. Set to <code>true</code> or <code>false</code> . |
| <code>mapreduce.marklogic.input.ssoptionclass</code> | The name of a class implementing <code>SslConfigOptions</code> , used to configure the SSL connection when <code>input.usessl</code> is <code>true</code> . |
| <code>mapreduce.marklogic.input.documentselector</code> | An XQuery path expression step specifying a sequence of document nodes to use for input. Basic mode only. Default: <code>fn:collection()</code> . See “Basic Input Mode” on page 40. |
| <code>mapreduce.marklogic.input.subdocumentexpr</code> | An XQuery path expression used with <code>input.documentselector</code> to select sub-document items to use in input key-value pairs. Basic mode only. See “Basic Input Mode” on page 40. |
| <code>mapreduce.marklogic.input.namespace</code> | A comma-separated list of namespace name-URI pairs to use when evaluating the path expression constructed from <code>input.documentselector</code> and <code>input.subdocumentexpr</code> . Basic mode only. See “Basic Input Mode” on page 40. |
| <code>mapreduce.marklogic.input.lexiconfunctionclass</code> | The class type of a lexicon function subclass. Use with <code>KeyValueInputFormat</code> and <code>ValueInputFormat</code> . Only usable in basic input mode. See “Using a Lexicon to Generate Key-Value Pairs” on page 41. |

| Property | Description |
|--|--|
| <code>mapreduce.marklogic.input.splitquery</code> | In advanced input mode, the query used to generate input splits. See “Creating Input Splits” on page 50. |
| <code>mapreduce.marklogic.input.query</code> | In advanced input mode, the query used to select input fragments from MarkLogic Server. See “Creating Input Key-Value Pairs” on page 52. |
| <code>mapreduce.marklogic.input.bindsplitrange</code> | Indicates whether or not your input query requires access to the external split range variables, <code>splitstart</code> and <code>splitend</code> . Only usable in advanced input mode. Default: <code>false</code> . See “Optimizing Your Input Query” on page 53. |
| <code>mapreduce.marklogic.input.maxsplitsize</code> | The maximum number of records (fragments) per input split. Default: 50,000. |
| <code>mapreduce.marklogic.input.keyclass</code> | The class type of the map phase input keys. Use with <code>KeyValueInputFormat</code> . See “Using <code>KeyValueInputFormat</code> and <code>ValueInputFormat</code> ” on page 56. |
| <code>mapreduce.marklogic.input.valueclass</code> | The class type of the map phase input values. Use with <code>KeyValueInputFormat</code> and <code>ValueInputFormat</code> . See “Using <code>KeyValueInputFormat</code> and <code>ValueInputFormat</code> ” on page 56. |
| <code>mapreduce.marklogic.input.recordtofragmentratio</code> | Defines the ratio of the number of retrieved input records to fragments. Used only for MapReduce’s progress reporting. Default: 1.0. |

4.7 InputFormat Subclasses

The MarkLogic Connector for Hadoop API provides subclasses of `InputFormat` for defining your map phase input splits and key-value pairs when using MarkLogic Server as an input source.

Specify the `InputFormat` subclass appropriate for your job using the

`org.apache.hadoop.mapreduce.job.setInputFormatClass` function. For example:

```
import com.marklogic.mapreduce.NodeInputFormat;
import org.apache.hadoop.mapreduce.Job;
...
public class LinkCountInDoc {
    ...
    public static void main(String[] args) throws Exception {
        Job job = new Job(conf);
        job.setInputFormatClass(NodeInputFormat.class);
    }
}
```

```

    }
    ...
}

```

The following table summarizes the `InputFormat` subclasses provided by the connector and the key and value types produced by each class. All classes referenced below are in the package `com.marklogic.mapreduce`. For more details, see the *MarkLogic Connector for Hadoop API Reference*.

| Class | Key Type | Value Type | Description |
|-----------------------------------|--------------------------|----------------------------|--|
| <code>MarkLogicInputFormat</code> | any | any | Abstract superclass for all connector specific <code>InputFormat</code> types. Generates input splits and key-value pairs from MarkLogic Server. |
| <code>DocumentInputFormat</code> | <code>DocumentURI</code> | <code>MarkLogicNode</code> | Generates key-value pairs using each node selected by the input query as a value and the URI of the document containing the node as a key. |
| <code>NodeInputFormat</code> | <code>NodePath</code> | <code>MarkLogicNode</code> | Generates key-value pairs using each node selected by the input query as a value and the node path of the node as a key. |

| Class | Key Type | Value Type | Description |
|---------------------|----------|------------|---|
| KeyValueInputFormat | any | any | Uses input from MarkLogic Server to generate input key-value pairs with user defined types. Use <code>mapreduce.marklogic.input.key-class</code> and <code>mapreduce.marklogic.input.valueclass</code> to specify the key and value types. See “Using KeyValueInputFormat and ValueInputFormat” on page 56. |
| ValueInputFormat | Int | any | Generates input key-value pairs where only the value is of interest. The key is simply a count of the number of items seen when the pair is generated. Use <code>mapreduce.marklogic.input.value-class</code> to specify the value type. See “Using KeyValueInputFormat and ValueInputFormat” on page 56. |

4.8 Lexicon Function Subclasses

The following table shows the correspondence between MarkLogic Server lexicon built-in functions and MarkLogic Connector for Hadoop lexicon function wrapper classes. Locate the lexicon function you wish to use and include a subclass of the corresponding wrapper class in your job. For details, see “Using a Lexicon to Generate Key-Value Pairs” on page 41.

| Lexicon Wrapper Superclass | XQuery Lexicon Function |
|----------------------------|---|
| ElemAttrValueCooccurrences | <code>cts:element-attribute-value-co-occurrences</code> |
| ElemValueCooccurrences | <code>cts:element-value-co-occurrences</code> |
| FieldValueCooccurrences | <code>cts:field-value-co-occurrences</code> |
| CollectionMatch | <code>cts:collection-match</code> |
| ElementAttributeValueMatch | <code>cts:element-attribute-value-match</code> |
| ElementAttributeWordMatch | <code>cts:element-attribute-word-match</code> |
| ElementValueMatch | <code>cts:element-value-match</code> |
| ElementWordMatch | <code>cts:element-word-match</code> |

| Lexicon Wrapper Superclass | XQuery Lexicon Function |
|----------------------------|---|
| FieldValueMatch | <code>cts:field-value-match</code> |
| FieldWordMatch | <code>cts:field-word-match</code> |
| UriMatch | <code>cts:uri-match</code> |
| WordMatch | <code>cts:word-match</code> |
| Collections | <code>cts:collections</code> |
| ElementAttributeValues | <code>cts:element-attribute-values</code> |
| ElementAttributeWords | <code>cts:element-attribute-words</code> |
| ElementValues | <code>cts:element-values</code> |
| ElementWords | <code>cts:element-words</code> |
| FieldValues | <code>cts:field-values</code> |
| FieldWords | <code>cts:field-words</code> |
| Uris | <code>cts:uris</code> |
| Words | <code>cts:words</code> |

5.0 Using MarkLogic Server for Output

This chapter covers the following topics related to storing MapReduce job results in a MarkLogic Server instance:

- [Basic Steps](#)
- [Creating a Custom Output Query with KeyValueOutputFormat](#)
- [Controlling Transaction Boundaries](#)
- [Performance Considerations for ContentOutputFormat](#)
- [Output Configuration Properties](#)
- [OutputFormat Subclasses](#)

5.1 Basic Steps

The MarkLogic Connector for Hadoop API supports storing MapReduce results in MarkLogic Server as documents, nodes, and properties. When using MarkLogic Server to store the results of the reduce phase, configuring the reduce step of a MapReduce job includes at least the major tasks:

- [Identifying the Output MarkLogic Server Instance](#)
- [Configuring the Output Key and Value Types](#)
- [Defining the Reduce Function](#)
- [Disabling Speculative Execution](#)
- [Example: Storing MapReduce Results as Nodes](#)

5.1.1 Identifying the Output MarkLogic Server Instance

The MarkLogic Server output instance is identified by setting configuration properties. For general information on setting configuration properties, see “Configuring a MapReduce Job” on page 32.

Specify the following properties to identify the output MarkLogic Server instance:

| Property | Description |
|--|---|
| <code>mapreduce.marklogic.output.host</code> | Hostname or IP address of the server hosting your output XDBC App Server. The host must be resolvable by the nodes in your Hadoop cluster, so you should usually not use “localhost”. |
| <code>mapreduce.marklogic.output.port</code> | The port configured for the target XDBC App Server on the input host. |
| <code>mapreduce.marklogic.output.username</code> | Username privileged to update the database attached to the XDBC App Server. |
| <code>mapreduce.marklogic.output.password</code> | Cleartext password for the <code>output.username</code> user. |

When you use MarkLogic Server in a cluster, all MarkLogic Server hosts containing a forest in the output database must be accessible through an XDBC server on the same port. The host identified in the configuration properties may be any qualifying host in the cluster. For details, see “Deploying the Connector with a MarkLogic Server Cluster” on page 11.

Note: The target database must be configured for manual directory creation by setting the `directory creation` database configuration setting to “manual”. See [Database Settings](#) in the *Administrator’s Guide*.

You can configure a job to connect to the App Server through SSL by setting the `mapreduce.marklogic.output.usessl` property. For details, see “Making a Secure Connection to MarkLogic Server with SSL” on page 13. For an example, see “ContentReader” on page 103.

For more information on the properties, see “Output Configuration Properties” on page 79.

5.1.2 Configuring the Output Key and Value Types

As discussed in “Reduce Task” on page 30, the `org.apache.hadoop.mapreduce.OutputFormat` subclass configured for the job determines the types of the output keys and values, and how results are stored. Use the `org.apache.hadoop.mapreduce.Job` API to configure the `OutputFormat` and the output key and value types for a job.

The Apache Hadoop MapReduce framework includes `OutputFormat` subclasses for saving results to files in HDFS. The MarkLogic Connector for Hadoop API includes `OutputFormat` subclasses for storing results as documents, nodes, or properties in a MarkLogic Server database. The output key and value types you configure must match the `OutputFormat` subclass. For example, `PropertyOutputFormat` expects `(DocumentURI, MarkLogicNode)` key-value pairs, so configure the job with `DocumentURI` as the key type and `MarkLogicNode` as the value type.

For a summary of the `OutputFormat` subclasses provided by the MarkLogic Connector for Hadoop, including the expected key and value types, see “OutputFormat Subclasses” on page 83.

The example below configures a job to use `NodeOutputFormat`, which expects `(NodePath, MarkLogicNode)` key-value pairs.

```
import com.marklogic.mapreduce.NodeOutputFormat;
import com.marklogic.mapreduce.NodePath;
import com.marklogic.mapreduce.MarkLogicNode;
import org.apache.hadoop.mapreduce.Job;
...
public class LinkCountInDoc {
    ...

    public static class IntSumReducer
        extends Reducer<Text, IntWritable, NodePath, MarkLogicNode> {...}
    public static void main(String[] args) throws Exception {
        Job job = new Job(conf);
        job.setOutputFormatClass(NodeOutputFormat.class);
        job.setOutputKeyClass(NodePath.class);
        job.setOutputValueClass(MarkLogicNode.class);
        job.setReducerClass(IntSumReducer.class);
        ...
    }
}
```

5.1.3 Defining the Reduce Function

To create a reducer, define a subclass of `org.apache.hadoop.mapreduce.Reducer` and override at least the `Reducer.reduce` method. Your `reduce` method should generate key-value pairs of the expected type and write them to the method’s `Context` parameter. The MapReduce framework subsequently stores the results written to the `Context` parameter in the file system or database using the configured `OutputFormat` subclass.

The output key and value types produced by the `reduce` method must match the output key and value types expected by the `OutputFormat` subclass. For example, if the job is configured to use `NodeOutputFormat` then the `reduce` method must generate `(NodePath, MarkLogicNode)` key-value pairs. The following example uses `(Text, IntWritable)` `reduce` input pairs and produces `(NodePath, MarkLogicNode)` output pairs using `NodeOutputFormat`:

```
import org.apache.hadoop.mapreduce.Reducer;
import org.apache.hadoop.io.Text;
import org.apache.hadoop.io.IntWritable;
import com.marklogic.mapreduce.MarkLogicNode;
import com.marklogic.mapreduce.NodePath;

public class LinkCountInDoc {
    ...
    public static class IntSumReducer
        extends Reducer<Text, IntWritable, NodePath, MarkLogicNode> {
```

```

    public void reduce(Text key, Iterable<IntWritable> values,
        Context context) throws IOException, InterruptedException {
        ...
        context.write(output_node_path, output_node);
    }
    ...
}

```

Notice that the `reduce` method receives a key and a list of all values produced by the map phase with the same key, as discussed in “MapReduce Overview” on page 25.

For a complete list of the `OutputFormat` subclasses provided by the connector and how they use the output keys and values to create database content, see “OutputFormat Subclasses” on page 83.

5.1.4 Disabling Speculative Execution

Disable Apache Hadoop MapReduce speculative execution when using MarkLogic Server for output.

Apache Hadoop MapReduce uses speculative execution to prevent jobs from stalling on slow tasks or worker nodes. For example, if most of the tasks for a job complete, but a few tasks are still running, Hadoop can schedule speculative redundant tasks on free worker nodes. If the original task completes before the redundant task, Hadoop cancels the redundant task. If the redundant task completes first, Hadoop cancels the original task.

Speculative execution is not safe when using MarkLogic Server for output because the cancelled tasks do not clean up their state. Uncommitted changes might be left dangling and eventually lead to `XDMP-FORESTTIM` errors.

Disable speculative execution by setting `mapred.map.tasks.speculative.execution` and/or `mapred.reduce.tasks.speculative.execution` to `false` in your job configuration file or using the `org.apache.conf.Configuration` API.

Set `mapred.map.tasks.speculative.execution` to `false` when using MarkLogic Server for output during map:

```

<property>
  <name>mapred.map.tasks.speculative.execution</name>
  <value>>false</value>
</property>

```

Set `mapred.reduce.tasks.speculative.execution` to `false` when using MarkLogic Server for output during map:

```

<property>
  <name>mapred.reduce.tasks.speculative.execution</name>

```

```
<value>false</value>
</property>
```

5.1.5 Example: Storing MapReduce Results as Nodes

This example demonstrates storing the results of a MapReduce job as child elements of documents in the database. The code samples shown here are small slices of the full code. For the complete code, see the `LinkCountInDoc` example in the `com.marklogic.mapreduce.examples` package.

The map step in this example creates a key-value pair for each href to a document in the collection. The reduce step sums up the number of references to each document and stores the total as a new `<ref-count>` element on the referenced document.

The output from the map phase is `(Text, IntWritable)` pairs, where the text is the title attribute text of an href and the integer is always 1, indicating one reference to the title. The sample content is such that the title in an href matches the leaf name of the containing document URI.

For example, if the sample data includes a document with the URI “`/space/wikipedia/enwiki/Drama film`” and three occurrences of the following href:

```
<a href="http://www.mediawiki.org/xml/export-0.4/" title="Drama film">
  drama film
</a>
```

Then the map output includes three key-value pairs of the form:

```
("Drama film", 1)
```

This results in an input key-value pair for reduce of the form:

```
("Drama film", (1, 1, 1))
```

The reduce phase computes the reference count (3, in this case) and attaches the count as a child node of the referenced document:

```
<wp:page>
  ...
  <ref-count>3</ref-count>
</wp:page>
```

To produce this result, the job uses `NodeOutputFormat` as the `OutputFormat` subclass. `NodeOutputFormat` expects the reducer to produce `(NodePath, MarkLogicNode)` output pairs. `NodeOutputFormat` uses the config property `com.marklogic.mapreduce.output.node.optype` to determine where to insert the node relative to the path. In this example, `nodeoptype` is set to `INSERT_CHILD`, so the `ref-count` node is inserted as a child of the node in the node path:

```
<property>
  <name>mapreduce.marklogic.output.node.optype</name>
```

```
<value>INSERT_CHILD</value>
</property>
```

To avoid the overhead of creating a new `MarkLogicNode` for every invocation of the `reduce` method, the sample overrides `Reducer.setup` to create a skeleton node which has its value replaced by the real reference count in `reduce`. The following code snippet demonstrates the initialization of the `result` and `element` variables used in `reduce`. For the full code, see the sample code for `com.marklogic.mapreduce.examples.LinkCountInDoc`.

```
import org.apache.hadoop.mapreduce.Reducer;
import org.apache.hadoop.io.IntWritable;
import org.apache.hadoop.io.Text;
import com.marklogic.mapreduce.MarkLogicNode;
import com.marklogic.mapreduce.NodeInputFormat;
import com.marklogic.mapreduce.NodeOutputFormat;
import com.marklogic.mapreduce.NodePath;
import javax.xml.parsers.DocumentBuilder;
import javax.xml.parsers.DocumentBuilderFactory;
import org.w3c.dom.Element;

public static class IntSumReducer
extends Reducer<Text, IntWritable, NodePath, MarkLogicNode> {
    private final static String TEMPLATE = "<ref-count>0</ref-count>";

    private Element element;
    private MarkLogicNode result;

    protected void setup(Context context)
        throws IOException, InterruptedException {
        try {
            DocumentBuilder docBuilder =
DocumentBuilderFactory.newInstance().newDocumentBuilder();
            InputStream sbis = new StringBufferInputStream(TEMPLATE);
            element = docBuilder.parse(sbis).getDocumentElement();
            result = new MarkLogicNode(element);
        }
        ...
    }
}
```

The following code snippet shows how the reducer sums up the input values, sets up the `ref-count` node, and builds the document node path from a known base URI and the href title passed in as key:

```
public static class IntSumReducer
extends Reducer<Text, IntWritable, NodePath, MarkLogicNode> {

    private final static String ROOT_ELEMENT_NAME = "//wp:page";
    private final static String BASE_URI_PARAM_NAME =
        "/space/wikipedia/enwiki/";
    private NodePath nodePath = new NodePath();
```

```

public void reduce(Text key, Iterable<IntWritable> values,
    Context context
    ) throws IOException, InterruptedException {
    ...

    // Compute reference count and store it in result node
    int sum = 0;
    for (IntWritable val : values) {
        sum += val.get();
    }
    element.setTextContent(Integer.toString(sum));

    // Build node path to referenced document
    StringBuilder buf = new StringBuilder();
    buf.append(BASE_URI).append(key);
    nodePath.setDocumentUri(buf.toString());
    nodePath.setRelativePath(ROOT_ELEMENT_NAME);

    // Store the final results for insertion into the database
    context.write(nodePath, result);
}

```

The above code, with an input pair of ("Drama film", (1,1,1)) produces the following (node path, node) output pair:

```

(
  /space/wikipedia/enwiki/Drama film//wp:page,
  <ref-count>3</ref-count>
)

```

Since the `mapreduce.marklogic.output.node.optype` property is `INSERT_CHILD`, the new `<ref-count>` node is inserted as a child of the `wp:page` element addressed by the node path.

With very little code modification, the sample can store the reference count as a property of the referenced document instead of as a child node. To do so, use `PropertyOutputFormat`, construct the document URI instead of node path, and set the config property `com.marklogic.mapreduce.output.propertyopttype`. For an example, see “[LinkCountInProperty](#)” on page 98.

5.2 Creating a Custom Output Query with `KeyValueOutputFormat`

Use `KeyValueOutputFormat` to create a custom output query that can manipulate arbitrary key and value data types and perform operations beyond updating documents, nodes, and properties. The topics covered by this section are:

- [Output Query Requirements](#)
- [Job Configuration](#)
- [Supported Type Transformations](#)

5.2.1 Output Query Requirements

A custom output query is a fully formed XQuery module, specified as the value of the configuration property `mapreduce.marklogic.output.query`. The MarkLogic Connector for Hadoop invokes the query for each output key-value pair produced by the map or reduce function configured to use `KeyValueOutputFormat`. Your query may perform any operations. Output returned by the query is discarded.

The output key-value pair available to your output query through external variables with the names “key” and “value”, in the namespace “`http://marklogic.com/hadoop`”. For example, the following output query prolog assumes a key type of `xs:string` and a value type of `element()`:

```
<property>
  <name>mapreduce.marklogic.output.query</name>
  <value><![CDATA[
    xquery version '1.0-ml';
    declare namespace mlmr = "http://marklogic.com/hadoop";
    declare variable $mlmr:key as xs:string external;
    declare variable $mlmr:value as element() external;
    (: use the key and value... :)
  ]]></value>
</property>
```

The key and value types are determined by the configuration properties `mapreduce.marklogic.output.keytype` and `mapreduce.marklogic.output.valuetype`. The key and value XQuery types are constrained to those XML schema types with a meaningful `org.apache.hadoop.io.Writable` type transformation, such as between `xs:string` and `org.apache.hadoop.io.Text` or between `element()` and `com.marklogic.mapreduce.MarkLogicNode`. For details, see “Supported Type Transformations” on page 75.

The prolog of your output query must include:

1. A namespace declaration for the namespace containing the key and value variables (“`http://marklogic.com/hadoop`”). You may use any namespace prefix.
2. An external variable declaration for `key` in the namespace from Step 1. Choose one of the types listed in “Supported Type Transformations” on page 75.
3. An external variable declaration for `value` in the namespace from Step 1. Choose one of the types listed in “Supported Type Transformations” on page 75.

Configure your job to use `KeyValueOutputFormat` and key and value Java types corresponding to the XQuery types expected by your output query. For details, see “Job Configuration” on page 73.

5.2.2 Job Configuration

Configure the following job properties to use a custom output query:

- Set `mapreduce.marklogic.output.query` to your output query. This must be a fully formed XQuery module, suitable for evaluation with `xmmp:eval`. For details, see “Output Query Requirements” on page 73.
- Set `mapreduce.marklogic.output.keytype` to the XQuery type of the key. The default type is `xs:string`.
- Set `mapreduce.marklogic.output.valuetype` to the XQuery type of the value. The default type is `xs:string`.
- Set the map or reduce output format to `KeyValueOutputFormat`.
- Set the map or reduce output key Java type to an `org.apache.hadoop.io.Writable` subclass that is convertible to the XQuery type in `mapreduce.marklogic.output.keytype`.
- Set the map or reduce output value Java type to an `org.apache.hadoop.io.Writable` that is subclass convertible to the XQuery type `mapreduce.marklogic.output.valuetype`.

For details on the available key and value types and supported conversions between the XQuery types and the Hadoop Java classes, see “Supported Type Transformations” on page 75.

The configuration properties may be set either in a configuration file or using the `org.apache.hadoop.mapreduce.Job` API. The following example configures the MarkLogic Connector for Hadoop output query, XQuery key type, and XQuery value type in a configuration file:

```
<property>
  <name>mapreduce.marklogic.output.keytype</name>
  <value>xs:string</value>
</property>
<property>
  <name>mapreduce.marklogic.output.valuetype</name>
  <value>element()</value>
</property>
<property>
  <name>mapreduce.marklogic.output.query</name>
  <value><![CDATA[
    xquery version '1.0-ml';
    declare namespace mlmr = "http://marklogic.com/hadoop";
    declare variable $mlmr:key as xs:string external;
    declare variable $mlmr:value as element() external;
    (: use the key and value... :)
  ]]></value>
</property>
```

The following example configures the job output format class, key class, and value class corresponding to the custom output query settings above:

```
import org.apache.hadoop.Job;
import org.apache.hadoop.conf.Configuration;
import org.apache.hadoop.io.Text;
import com.marklogic.mapreduce.MarkLogicNode;
```

```
class myClass {
    public static void main(String[] args) throws Exception {
        Configuration conf = new Configuration();
        Job job = new Job(conf);
        job.setOutputFormatClass(KeyValueOutputFormat);
        job.setOutputKeyClass(Text.class);
        job.setOutputValueClass(MarkLogicNode.class);
        ...
    }
}
```

For a complete example, see “RevisionGrouper” on page 101 and the sample code for `com.marklogic.mapreduce.examples.RevisionGrouper`.

5.2.3 Supported Type Transformations

When you use `KeyValueOutputFormat`, the MarkLogic Connector for Hadoop converts the key and value produced by the job map or reduce function from an `org.apache.hadoop.io.Writable` object to values of the configured XQuery types. The table below summarizes the supported type conversions.

When using `org.apache.hadoop.io.Text` with any XQuery type other than `xs:string`, the string value format must correspond to the serialized representation of the XQuery type. For example, if the XQuery type is `cts:box`, then the `Text` string representation must be a serialized `cts:box` value, such as “[`-122, 78, 30, 45`]”. Similarly, if the XQuery type is `xs:duration`, then the `Text` string representation must be a serialized `xs:duration`, such as “`-P92M`”.

| If the key or value XQuery type is | Then the Java type must be |
|--|--|
| xs:string | org.apache.hadoop.io.Text |
| xs:boolean | org.apache.hadoop.io.BooleanWritable org.apache.hadoop.io.Text |
| xs:integer | org.apache.hadoop.io.IntWritable org.apache.hadoop.io.VIntWritable org.apache.hadoop.io.LongWritable org.apache.hadoop.io.VLongWritable org.apache.hadoop.io.Text |
| xs:decimal | org.apache.hadoop.io.IntWritable org.apache.hadoop.io.VIntWritable org.apache.hadoop.io.LongWritable org.apache.hadoop.io.VLongWritable org.apache.hadoop.io.FloatWritable org.apache.hadoop.io.DoubleWritable org.apache.hadoop.io.Text |
| xs:float xs:double | org.apache.hadoop.io.IntWritable org.apache.hadoop.io.VIntWritable org.apache.hadoop.io.LongWritable org.apache.hadoop.io.VLongWritable org.apache.hadoop.io.FloatWritable org.apache.hadoop.io.DoubleWritable org.apache.hadoop.io.Text |
| xs:duration xs:dayTimeDuration xs:yearMonthDuration xs:dateTime xs:time xs:date | org.apache.hadoop.io.Text |
| xs:hexBinary | org.apache.hadoop.io.BytesWritable org.apache.hadoop.io.Text |
| xs:base64Binary | org.apache.hadoop.io.Text |
| xs:gDay xs:gMonth xs:gYear xs:gYearMonth | org.apache.hadoop.io.Text |
| cts:box cts:circle cts:point cts:polygon | org.apache.hadoop.io.Text |
| node() element() binary() | com.marklogic.mapreduce.MarkLogicNode org.apache.hadoop.io.Text |

5.3 Controlling Transaction Boundaries

When you use one of the `OutputFormat` subclasses provided by the MarkLogic Connector for Hadoop, such as `ContentOutputFormat`, the MarkLogic Connector for Hadoop manages the output MarkLogic Server XDBC session and the requests for storing results in the database.

Each output key-value pair represents an update to the database. Multiple updates are grouped into a single transaction. For all `OutputFormat` subclasses except `ContentOutputFormat`:

- Each output key-value pair generates an update request to MarkLogic Server.
- Every 1000 requests constitutes a transaction.

Use the `mapreduce.marklogic.output.transactionsize` configuration property to adjust the number of requests per transaction. For example, if you set the transaction size to 5, then the connector commits a transaction for every 5 requests. For `OutputFormat` classes other than `ContentOutputFormat`, this means the connector bundles every 5 updates into a transaction.

```
<property>
  <name>mapreduce.marklogic.output.transactionsize</name>
  <value>5</value>
</property>
```

For `ContentOutputFormat`, the interactions between the number of updates, requests, and transactions is more complicated. For details, see “Time vs. Space: Configuring Batch and Transaction Size” on page 77.

5.4 Performance Considerations for ContentOutputFormat

When using `ContentOutputFormat`, be aware of the following performance tradeoffs discussed in this section:

- [Time vs. Space: Configuring Batch and Transaction Size](#)
- [Time vs. Correctness: Using Direct Forest Updates](#)

5.4.1 Time vs. Space: Configuring Batch and Transaction Size

When using `ContentOutputFormat`, you can tune the insertion throughput and memory requirements of your job by configuring the batch size and transaction size of the job:

- `mapreduce.marklogic.output.batchsize` controls the number of output records (updates) per request to the server.
- `mapreduce.marklogic.output.transactionsize` controls the number of requests to the server per transaction.

Selecting a batch size is a speed vs. memory tradeoff. Each request to the server introduces overhead because extra work must be done. However, all the updates in a batch stay in memory until a request is sent, so larger batches consume more memory.

Transactions introduce overhead on MarkLogic Server, so performing multiple updates per transaction can improve insertion throughput. However, an open transaction holds locks on fragments with pending updates, potentially increasing lock contention and affecting overall application performance.

The default batch size is 100. If you do not explicitly set the transaction size, it varies with batch size, adjusting to keep the maximum number of updates per transaction at 2000.

Consider the following example of inserting 10000 documents. Assume batch size is controlled by setting `mapreduce.marklogic.output.batchsize` and transaction size not explicitly set. The last row represents the default behavior for `ContentOutputFormat`.

| batch size | total requests | txn size | total txns | updates/txn |
|------------|----------------|----------|------------|-------------|
| 1 | 10000 | 2000 | 5 | 2000 |
| 10 | 1000 | 200 | 5 | 2000 |
| 100 | 100 | 20 | 5 | 2000 |

If you explicitly set `mapreduce.marklogic.output.transactionsize`, then transaction size does not vary based on batch size. Given the example above of inserting 10000 documents, if you explicitly set batch size to 100 and transaction size to 50, then the job requires 2 transactions and each transaction performs 5000 updates.

Batch size is not configurable for other output formats, such as `NodeOutputFormat`, `PropertyOutputFormat`, and `KeyValueFormat`. For these classes, batch size is always 1 and the default transaction size is 1000. See “Controlling Transaction Boundaries” on page 77.

5.4.2 Time vs. Correctness: Using Direct Forest Updates

`ContentOutputFormat` performs best when the updating tasks interact directly with the forests in which content is inserted. However, direct forest updates can create duplicate document URIs under the following circumstances:

- Content with the same URI already exists in the database, and
- The content was inserted using forest placement or the number of forests changed after initial document creation.

Forest placement occurs when you use the `$forest-ids` parameter of `xdmp:document-insert` to instruct MarkLogic Server to insert a document in a specific forest or to choose from a specific set of forests. See [Specifying a Forest in Which to Load a Document](#) in the *Application Developer’s Guide*.

To prevent duplicate URIs, the MarkLogic Connector for Hadoop defaults to a slower protocol for `ContentOutputFormat` when it detects the potential for updates to existing content. In this case, MarkLogic Server manages the forest selection, rather than the MarkLogic Connector for Hadoop. This behavior guarantees unique URIs at the cost of performance.

You may override this behavior and use direct forest updates by doing the following:

- Set `mapreduce.marklogic.output.content.directory`. This guarantees all inserts will be new documents. If the output directory already exists, it will either be removed or cause an error, depending on the value of `mapreduce.marklogic.output.content.cleandir`.
- Set `mapreduce.marklogic.output.content.fastload` to `true`. When `fastload` is `true`, the MarkLogic Connector for Hadoop always optimizes for performance, even if duplicate URIs are possible.

You can safely set `mapreduce.marklogic.output.content.fastload` to `true` if the number of forests in the database will not change while the job runs, and at least one of the following is true:

- Your job only creates new documents. That is, you are certain that the URIs are not in use by any document or property fragments already in the database.
- The URIs output with `ContentOutputFormat` may already be in use, but both these conditions are true:
 - The in-use URIs were not originally inserted using forest placement.
 - The number of forests in the database has not changed since initial insertion.
- You set `mapreduce.marklogic.output.content.directory`.

5.5 Output Configuration Properties

The table below summarizes connector configuration properties for using MarkLogic Server as an output destination. For details, see `com.marklogic.mapreduce.MarkLogicConstants` in the *MarkLogic Connector for Hadoop API Reference*.

| Property | Description |
|---|---|
| <code>mapreduce.marklogic.output.username</code> | Username for a user privileged to write to the database attached to your XDBC App Server. |
| <code>mapreduce.marklogic.output.password</code> | Cleartext password for the <code>output.username</code> user. |
| <code>mapreduce.marklogic.output.host</code> | Hostname of the server hosting your output XDBC App Server. |
| <code>mapreduce.marklogic.output.port</code> | The port configured for the XDBC App Server on the output host. |
| <code>mapreduce.marklogic.output.usessl</code> | Whether or not to use an SSL connection to the server. Default: <code>false</code> . |
| <code>mapreduce.marklogic.output.ssloptionsclass</code> | The name of a class implementing <code>com.marklogic.mapreduce.SslConfigOptions</code> . Used to configure the SSL connection when <code>output.usessl</code> is true. |
| <code>mapreduce.marklogic.output.node.namespace</code> | A comma-separated list of namespace name-URI pairs to use when evaluating element names in the node path with <code>NodeOutputFormat</code> . |
| <code>mapreduce.marklogic.output.batchsize</code> | The number of output key-value pairs (updates) to send to MarkLogic Server in a single request. Only honored by <code>ContentOutputFormat</code> . Default: 100. Other output formats have an implicit, unconfigurable batch size of 1. |
| <code>mapreduce.marklogic.output.content.type</code> | When using <code>ContentOutputFormat</code> , specifies the content type of output documents. Set to one of <code>XML</code> , <code>TEXT</code> , or <code>BINARY</code> . |

| Property | Description |
|---|--|
| mapreduce.marklogic.output.content.directory | The database directory in which to create output documents when using <code>ContentOutputFormat</code> . If <code>content.cleandir</code> is false (the default), then the directory must not already exist. If <code>content.cleandir</code> is true and the directory already exists, it is deleted as part of job submission. |
| mapreduce.marklogic.output.content.collection | A comma separated list of collections to which output documents are added when using <code>ContentOutputFormat</code> . |
| mapreduce.marklogic.output.content.permission | A comma separated list of role-capability pairs to associate with output documents when using <code>ContentOutputFormat</code> . |
| mapreduce.marklogic.output.content.quality | The document quality to use when creating output documents with <code>ContentOutputFormat</code> . Default: 0. |
| mapreduce.marklogic.output.keytype | The XQuery type of the output key available to the query defined by <code>mapreduce.marklogic.output.query</code> . Default: <code>xs:string</code> . |
| mapreduce.marklogic.output.valuetype | The XQuery type of the output value available to the query defined by <code>mapreduce.marklogic.output.query</code> . Default: <code>xs:string</code> . |
| mapreduce.marklogic.output.query | A custom output query to be used by <code>KeyValueOutputFormat</code> . See “Creating a Custom Output Query with <code>KeyValueOutputFormat</code> ” on page 72. |

| Property | Description |
|---|--|
| <code>mapreduce.marklogic.output.content.cleandir</code> | Whether or not to remove the database directory specified by <code>output.content.directory</code> before storing the reduce phase results. If false, an error occurs if the directory already exists when the job runs. Default: false. |
| <code>mapreduce.marklogic.output.content.fastload</code> | Whether or not to force optimal performance when using <code>ContentOutputFormat</code> . Setting this property to true can result in duplicate URIs. Default: false. See “Performance Considerations for <code>ContentOutputFormat</code> ” on page 77. |
| <code>mapreduce.marklogic.output.node.optype</code> | When using <code>NodeOutputFormat</code> , the node operation to perform. Set to one of: <code>REPLACE</code> , <code>INSERT_BEFORE</code> , <code>INSERT_AFTER</code> , <code>INSERT_CHILD</code> . |
| <code>mapreduce.marklogic.output.property.optype</code> | When using <code>PropertyOutputFormat</code> , the property operation to perform. Set to one of <code>SET_PROPERTY</code> or <code>ADD_PROPERTY</code> . |
| <code>mapreduce.marklogic.output.property.alwayscreate</code> | When using <code>PropertyOutputFormat</code> , whether or not to create a property even if no document exists with the output document URI. Default: false. |
| <code>mapreduce.marklogic.output.transactionsize</code> | The number of requests to MarkLogic Server per transaction. Default: For <code>ContentOutputFormat</code> , this varies with <code>mapreduce.marklogic.output.batchsize</code> to maintain 2000 updates/transaction; for other output formats, 1000. |

5.6 OutputFormat Subclasses

The MarkLogic Connector for Hadoop API provides subclasses of `OutputFormat` for defining your reduce output key-value pairs and storing the results in a MarkLogic Server database.

Specify the `OutputFormat` subclass appropriate for your job using the

`org.apache.hadoop.mapreduce.job.setOutputFormatClass` function. For example:

```
import com.marklogic.mapreduce.NodeInputFormat;
import org.apache.hadoop.mapreduce.Job;
...
public class LinkCountInDoc {
    ...
    public static void main(String[] args) throws Exception {
        Job job = new Job(conf);
        job.setInputFormatClass(NodeOutputFormat.class);
        ...
    }
}
```

The following table summarizes the `OutputFormat` subclasses provided by the MarkLogic Connector for Hadoop and the key and value types produced by each class. All classes referenced below are in the package `com.marklogic.mapreduce`. All referenced properties are covered in “Output Configuration Properties” on page 79. For more details on these classes and properties, see the *MarkLogic Connector for Hadoop API Reference*.

| Class | Key Type | Value Type | Description |
|------------------------------------|--------------------------|-------------------------|---|
| <code>MarkLogicOutputFormat</code> | any | any | Superclass for all connector-specific <code>OutputFormat</code> classes. Stores output in a MarkLogic Server database. |
| <code>ContentOutputFormat</code> | <code>DocumentURI</code> | any (text, XML, binary) | Stores output in a MarkLogic Server database, using the key as the document URI and the value as the content. The content type is determined by the <code>mapreduce.marklogic.output.content.type</code> config property. Related configuration properties: <ul style="list-style-type: none"> <code>content.type</code> <code>content.directory</code> <code>content.collection</code> <code>content.permission</code> <code>content.quality</code> |

| Class | Key Type | Value Type | Description |
|----------------------|-------------|---------------|---|
| NodeOutputFormat | NodePath | MarkLogicNode | Stores output in a MarkLogic Server database, using the key as the node path and the value as the node to insert. The <code>mapreduce.marklogic.output.nodeopt-type</code> config property controls where the node is inserted relative to the node path in the key. |
| PropertyOutputFormat | DocumentURI | MarkLogicNode | Stores output in a MarkLogic Server database by inserting the value node as a property of the document in the key URI. The <code>mapreduce.marklogic.output.property.optype</code> config property controls how the property is set. |
| KeyValueOutputFormat | any | any | Run the query defined by <code>mapreduce.marklogic.output.query</code> with each output key-value pair. The result is determined by the output query. The key and value types are determined by <code>mapreduce.marklogic.output.keytype</code> and <code>mapreduce.marklogic.output.valuetype</code> . See “Creating a Custom Output Query with Key-ValueOutputFormat” on page 72. |

6.0 Troubleshooting and Debugging

This chapter covers the following topics related to troubleshooting and debugging MapReduce jobs that use the MarkLogic Connector for Hadoop:

- [Enabling Debug Level Logging](#)
- [Solutions to Common Problems](#)

6.1 Enabling Debug Level Logging

Enable debug logging for more insight into the job processing. For example, with debug logging enabled, the MarkLogic Connector for Hadoop logs detailed information about input splits, including the split query. To enable debug logging:

1. Edit `$HADOOP_HOME/conf/log4j.properties`.
2. Add the following line to enable debug logging:

```
log4j.logger.com.marklogic.mapreduce=DEBUG
```

3. Re-run the job and look for `DEBUG` messages in the console output or log files.

In standalone mode, job output goes to `stdout` and `stderr`. In pseudo-distributed or fully distributed mode, job output goes to `stdout`, `stderr`, and Hadoop log files. For information on locating log files, see “Viewing Job Status and Logs” on page 34.

6.2 Solutions to Common Problems

This section presents symptoms and possible solutions for the following common problems:

- [Configuration File Not Found](#)
- [XDBC App Server Not Reachable](#)
- [Authorization Failure](#)

See the Apache Hadoop wiki for additional tips:

<http://wiki.apache.org/hadoop/TroubleShooting>

6.2.1 Configuration File Not Found

| | |
|-----------|---|
| Symptom: | Job exits with <code>IllegalArgumentException</code> related to one of the connector configuration properties. |
| Cause: | Apache Hadoop failed to find or open the sample configuration file. |
| Solution: | Confirm that job configuration file is present in <code>\$HADOOP_HOME/conf</code> and is readable by the user under which Apache Hadoop runs. |
| Example: | Exception in thread "main" java.lang.IllegalArgumentException: mapreduce.marklogic.output.hosts is not specified |

6.2.2 XDBC App Server Not Reachable

| | |
|-----------|--|
| Symptom: | Job exits with an error related to the default provider, or exits with a connection refused error. |
| Cause: | The XDBC App Server is not reachable. Either the host or port for the input or output XDBC App Server is incorrect. |
| Solution: | Correct your configuration settings in the job configuration file in <code>\$HADOOP_HOME/conf</code> . |
| Example: | Exception in thread "main" java.lang.IllegalArgumentException: Default provider - Not a usable net address Exception in thread "main" java.io.IOException: com.marklogic.xcc.exceptions.ServerConnectionException: Connection refused |

6.2.3 Authorization Failure

| | |
|-----------|---|
| Symptom: | Job exits with an authorization failure. |
| Cause: | The user name or password for the input or output user is incorrect. |
| Solution: | Correct your configuration settings in the job configuration file in <code>\$HADOOP_HOME/conf</code> . |
| Example: | Exception in thread "main" java.io.IOException: com.marklogic.xcc.exceptions.RequestPermissionException: Authorization failed for user 'fred' |

7.0 Using the Sample Applications

This chapter covers the following topics related to using the sample applications:

- [Set Up for All Samples](#)
- [Additional Sample Data Setup](#)
- [Interacting with HDFS](#)
- [Sample Applications](#)

7.1 Set Up for All Samples

The following topics apply to preparing to run all the sample applications:

1. [Install Required Software](#)
2. [Install the Sample Configuration Files](#)
3. [Modify the Sample Configuration Files](#)

The `LinkCount` samples, such as `LinkCountInDoc` and `LinkCountValue`, require additional preparation. See “Additional Sample Data Setup” on page 90.

For details about the individual samples, see “Sample Applications” on page 95.

7.1.1 Install Required Software

Install and configure MarkLogic Server, Apache Hadoop MapReduce, and the MarkLogic Connector for Hadoop. For instructions, see “Getting Started with the MarkLogic Connector for Hadoop” on page 15.

Note: The MarkLogic Connector for Hadoop only supports Apache Hadoop MapReduce version 0.20.2.

The samples require at least one MarkLogic Server database and XDBC App Server. The `LinkCount` samples require a specific database configuration and data set; see “Additional Sample Data Setup” on page 90. The other samples may be run against any XDBC App Server and database, including the `hello-hadoop-xdbc` App Server and `hello-hadoop` database created in [Getting Started with the MarkLogic Connector for Hadoop](#).

7.1.1.1 Multi-host Configuration Considerations

“Getting Started with the MarkLogic Connector for Hadoop” on page 15 describes setting up a single-host configuration, where MarkLogic Server, Apache Hadoop MapReduce, and the MarkLogic Connector for Hadoop are installed on the same host, and Apache Hadoop MapReduce is configured for standalone operation. A multi-host configuration, with Apache Hadoop MapReduce configured for pseudo-distributed or fully-distributed operation, more accurately represents a production deployment.

If you choose to use a multi-host, distributed configuration be aware of the following:

- The MarkLogic Server host configured for the job must be reachable by hostname from the Apache Hadoop MapReduce worker nodes.
- The MarkLogic Connector for Hadoop must be installed on the Apache Hadoop MapReduce host on which you run the sample jobs.
- Normally, you can use different MarkLogic Server instances for input and output, but the `LinkCount` samples expect the same database for both input and output.

Some of the samples use Apache Hadoop HDFS for input or output. If Hadoop is configured for pseudo- or fully-distributed operation, HDFS must be initialized before running the samples.

To check whether or not HDFS is initialized, run the following command. It should run without error and show similar output:

```
$ hadoop dfs -ls /
drwxr-xr-x - marklogic\me mygroup 0 2011-07-19 10:48 /tmp
drwxr-xr-x - marklogic\me mygroup 0 2011-07-19 10:51 /user
```

If the command fails, HDFS might not be initialized. See “Initializing HDFS” on page 93.

7.1.2 Install the Sample Configuration Files

The sample applications include MapReduce configuration files containing MarkLogic Connector for Hadoop settings. These configuration files must be located in the `conf/` directory of the Hadoop MapReduce installation on the Hadoop Job Tracker host. To learn about the Job Tracker, see “Understanding the MapReduce Job Life Cycle” on page 28.

Copy the configuration files in `$CONNECTOR_HOME/conf` to `$HADOOP_HOME/conf`. For example, if you install the MarkLogic Connector for Hadoop to `/space/marklogic-mapreduce` and your Apache Hadoop MapReduce installation is in `/space/hadoop`, then copy the configuration files as follows:

```
$ cp /space/marklogic-mapreduce/conf/*.xml /space/hadoop/conf
```

7.1.3 Modify the Sample Configuration Files

For each sample you plan to run, modify the MarkLogic Connector for Hadoop sample configuration file in `$HADOOP_HOME/conf/` to match your MarkLogic Server configuration. The configuration file associated with each sample is listed below.

| Sample | Configuration File |
|--|------------------------------|
| HelloWorld | marklogic-hello-world.xml |
| LinkCountInDoc | marklogic-nodein-nodeout.xml |
| LinkCountInProperty | marklogic-textin-propout.xml |
| LinkCountValue | marklogic-textin-textout.xml |
| LinkCountCooccurrences | marklogic-lexicon.xml |
| LinkCount | marklogic-advanced.xml |
| RevisionGrouper | marklogic-nodein-qryout.xml |
| BinaryReader | marklogic-subbinary.xml |
| ContentReader | marklogic-docin-textout.xml |
| ContentLoader | marklogic-textin-docout.xml |
| ZipContentLoader | marklogic-textin-docout.xml |

The configuration properties requiring modification vary from sample to sample. For example, a sample which uses MarkLogic Server for input and HDFS for output will not include `mapreduce.marklogic.output.*` properties.

If the sample uses MarkLogic Server for input, modify at least the following config properties. For details, see “Identifying the Input MarkLogic Server Instance” on page 36.

- `mapreduce.marklogic.input.username`
- `mapreduce.marklogic.input.password`
- `mapreduce.marklogic.input.host`
- `mapreduce.marklogic.input.port`

If the sample uses MarkLogic Server for output, modify at least the following config properties. For details, see “Identifying the Output MarkLogic Server Instance” on page 66.

- `mapreduce.marklogic.output.username`
- `mapreduce.marklogic.output.password`
- `mapreduce.marklogic.output.host`
- `mapreduce.marklogic.output.port`

Some samples might require additional customization. For details on a specific sample, see “Sample Applications” on page 95.

7.2 Additional Sample Data Setup

The following samples require a special database configuration and input data set: The `LinkCount*` samples (`LinkCountInDoc`, `LinkCountValue`, etc.), `RevisionGroup`. If you do not plan to run these samples, you may skip this section.

This section walks you through creating the MarkLogic Server environment required by these samples.

- [Creating the Database](#)
- [Creating the XDBC App Server](#)
- [Loading the Data](#)

7.2.1 Creating the Database

Use the following information to create a database named “hadoop-samples” with 2 forests and 2 attribute range indexes. You may use a different database name. Use the defaults for any configuration parameters not mentioned below.

For detailed instructions, see [Creating and Configuring Forests and Databases](#) and [Defining Attribute Range Indexes](#) in the *Administrator’s Guide*.

| Configuration Parameter | | Setting |
|-------------------------|----------------------|--|
| database name | | hadoop-samples |
| forest names | | hadoop-samples-1, hadoop-samples-2 |
| attribute range index 1 | scalar type | string |
| | parent namespace uri | http://www.mediawiki.org/xml/export-0.4/ |
| | parent localname | a |
| | localname | href |
| | collation | Unicode Codepoint |
| range value positions | | true |

| Configuration Parameter | | Setting |
|----------------------------|----------------------|--|
| attribute range index 2 | scalar type | string |
| | parent namespace uri | http://www.mediawiki.org/xml/export-0.4/ |
| | parent localname | a |
| | localname | title |
| | collation | Unicode Codepoint |
| range value positions | | true |

7.2.2 Creating the XDBC App Server

Use the following information to create an XDBC App Server and attach it to the database created in the previous section. You may use a different name and port. The settings below are the ones assumed by this guide.

For detailed instructions, see [Creating and Configuring App Servers](#) in the *Administrator's Guide*.

| Configuration Parameter | Setting |
|-------------------------|---------------------|
| xdbc server name | hadoop-samples-xdbc |
| root | (any) |
| port | 9002 |
| database | hadoop-samples |

7.2.3 Loading the Data

Load the data from `$(CONNECTOR_HOME)/sample-data` into the `hadoop-samples` database with a URI prefix of `enwiki/`. The query below is an example of loading the data using InfoStudio, assuming `$(CONNECTOR_HOME)` is `/space/marklogic/mapreduce`. To use this query, modify the value of `$(local-dir)` to match your environment, then run the query in Query Console against any Content Source. To learn more about Query Console, see [Query Console Walkthrough](#) in the *Query Console User Guide*.

```
(: Load the sample data into hadoop-samples database.
: Modify $(local-dir) to match your environment. Run against any
: content source. :)

xquery version "1.0-ml";
```

```

import module namespace info =
"http://marklogic.com/appservices/infostudio"
  at "/MarkLogic/appservices/infostudio/info.xqy";

(: Replace the value of $local-dir with your path to the sample data
directory :)
declare variable $local-dir :=
  "/space/marklogic/mapreduce/sample-data";

declare variable $samples-db := "hadoop-samples";

info:load(
  $local-dir, (),
  <options xmlns="http://marklogic.com/appservices/infostudio">
    <uri>
      <literal>enwiki</literal>
      <path strip-prefix="{ $local-dir }"/><literal>/</literal>
      <filename/>
    </uri>
    <format>xml</format>
  </options>,
  $samples-db)

```

The query reports an InfoStudio ticket number when it completes. Use the ticket information with `info:ticket-errors` to check for load errors. For example:

```

xquery version "1.0-m1";
declare namespace html = "http://www.w3.org/1999/xhtml";

import module namespace info =
"http://marklogic.com/appservices/infostudio"
  at "/MarkLogic/appservices/infostudio/info.xqy";

(: Replace the ticket number with your ticket :)
info:ticket-errors("/tickets/ticket/8667452332156256161")

```

If the load succeeds, `info:ticket-errors` should report zero errors in the `total` attribute, similar to the following:

```
<info:errors start="1" page-length="10" total="0" ... />
```

7.3 Interacting with HDFS

Some of the samples use Apache Hadoop HDFS for input or output. This section briefly summarizes how to copy data into or retrieve data from HDFS in when using Hadoop in pseudo-distributed or fully-distribute configurations.

If you use Apache Hadoop MapReduce standalone, you may skip this section. Standalone Hadoop is the configuration created in “Getting Started with the MarkLogic Connector for Hadoop” on page 15. In a standalone configuration, HDFS uses the local file system directly. You do not need to initialize HDFS, and you can use normal Unix commands to work with the input and output files. You may still use HDFS commands to examine the file system.

This section covers following topics related to pseudo- and fully-distributed HDFS operation:

- [Initializing HDFS](#)
- [Accessing Results Saved to HDFS](#)
- [Placing Content in HDFS to Use as Input](#)

Use the following command see all available HDFS commands:

```
$ hadoop dfs -help
```

For more details, see the Apache Hadoop HDFS documentation at <http://hadoop.apache.org/hdfs/>.

7.3.1 Initializing HDFS

If you use Apache Hadoop MapReduce in pseudo-distributed or fully-distributed mode, then you must initialize HDFS by formatting it. Run the following command to initialize HDFS:

```
$ $HADOOP_HOME/bin/hadoop namenode -format
```

Near the end of the output, you should see a message that HDFS has been successfully formatted:

```
...
*****/
11/10/03 09:35:14 INFO namenode.FSNamesystem:...
11/10/03 09:35:14 INFO namenode.FSNamesystem: supergroup=supergroup
11/10/03 09:35:14 INFO namenode.FSNamesystem: isPermissionEnabled=true
11/10/03 09:35:14 INFO common.Storage: Image file of size 98 saved ...
11/10/03 09:35:14 INFO common.Storage: Storage directory
/tmp/hadoop-sample/dfs/name has been successfully formatted.
11/10/03 09:35:14 INFO namenode.NameNode: SHUTDOWN_MSG:
/*****
SHUTDOWN_MSG: Shutting down NameNode at sample.marklogic.com...
*****/
```

If formatting succeeds, you can successfully use the ls command to examine HDFS:

```
$ $HADOOP_HOME/bin/hadoop dfs -ls /
```

For information about pseudo-distributed and fully-distributed operation, see the Apache Hadoop MapReduce documentation:

http://hadoop.apache.org/common/docs/current/mapred_tutorial.html

7.3.2 Accessing Results Saved to HDFS

Some of the sample applications store results to HDFS. You can browse HDFS and examine results from the command line or through your web browser.

In pseudo-distributed or fully distributed configurations, HDFS output pathnames given on the command line of an example are relative to `/user/your_username` in HDFS by default. For example if you run the `LinkCountValue` example, which saves results to HDFS, and specify the output directory as `linkcountvalue`, then the results are in HDFS under `/user/your_username/linkcountvalue`.

To access HDFS through your web browser, use the HDFS NameNode administration page. By default, this interface is available on port 50070 on the NameNode host. Assuming localhost is the NameNode, browse to this URL and click on the “Browse the file system link” near the top of the page to browse HDFS:

```
http://localhost:50070
```

To browse HDFS from the command line, use a command similar to the following:

```
$ hadoop dfs -ls /user/your_username
```

For example if you run the `LinkCountValue` example and specify the output directory as `linkcountvalue`, you would see results similar to the following, after running the example:

```
$ hadoop dfs -ls /user/me/linkcountvalue
drwxr-xr-x  - me mygroup ... /user/me/linkcountvalue/_logs
-rw-r--r--  1 me mygroup ... /user/me/linkcountvalue/part-r-00000
```

The results are in the `part-r-xxxxx` file. To see the last few lines of the results, use a command similar to the following:

```
$ hadoop dfs -tail /user/me/linkcountvalue/part-r-00000
```

To copy the result from HDFS to your system’s file system, use a command similar to the following:

```
$ hadoop dfs -get /user/me/linkcountvalue/part-r-00000 \
/my/destination/linkcountvalue.txt
```

7.3.3 Placing Content in HDFS to Use as Input

Some of the samples use HDFS for input. These samples require you to copy the input data to HDFS before running the sample. Place the input files under `/user/your_username` in HDFS using a command such as the following:

```
$ hadoop dfs -put ./mycontent.zip /user/me/zipcontentloader
```

Relative pathnames are relative to `/user/your_username`, so to check the file copied into HDFS above, use a command similar to the following:

```
$ hadoop dfs -ls /user/me/zipcontentloader
-rw-r--r--  1 me mygroup ... /user/me/zipcontentloader/mycontent.zip
```

When you copy files into HDFS, there must not be a pre-existing file of the same name.

7.4 Sample Applications

This section contains detailed instructions for running each of the samples summarized in the table below.

The MarkLogic Connector for Hadoop distribution includes the following resources related to the samples:

- Source code, in `$CONNECTOR_HOME/src`.
- Compiled code, in `$CONNECTOR_HOME/lib/marklogic-mapreduce-examples-version.jar`.
- Javadoc. See the package `com.marklogic.mapreduce.examples` in the Javadoc under `$CONNECTOR_HOME/docs`.

The sample applications are:

| Sample | Input | Output | Description |
|--|------------------|------------------|--|
| HelloWorld | MarkLogic Server | MarkLogic Server | Reads the first word from text in input XML documents, concatenates the words, then stores the results as a new text document in MarkLogic Server. |
| LinkCountInDoc | MarkLogic Server | MarkLogic Server | Counts href link title attributes in documents in MarkLogic Server, then stores the count as a child node of the referenced document. |
| LinkCountInProperty | MarkLogic Server | MarkLogic Server | Counts href link title attributes in documents in MarkLogic Server, then stores the count as a property of the referenced document. |
| LinkCountValue | MarkLogic Server | HDFS | Counts href link titles attributes in documents in MarkLogic Server, then stores the counts in HDFS text files. |
| LinkCountCooccurrences | MarkLogic Server | HDFS | Counts href link title attributes in documents in MarkLogic Server using a lexicon function, then stores the counts in HDFS text files. |
| LinkCount | MarkLogic Server | HDFS | Equivalent to <code>LinkCountValue</code> , but demonstrates using advanced input mode to provide your own input split and input queries. |

| Sample | Input | Output | Description |
|----------------------------------|------------------|------------------|---|
| RevisionGrouper | MarkLogic Server | MarkLogic Server | Demonstrates the use of a custom output query, using <code>KeyValueOutputFormat</code> . |
| BinaryReader | MarkLogic Server | HDFS | Demonstrates using advanced input mode with an input query optimized using the split range. |
| ContentReader | MarkLogic Server | HDFS | Reads documents in a MarkLogic Server database, using an SSL-enabled connection, then writes the contents to HDFS text files. |
| ContentLoader | HDFS | MarkLogic Server | Reads text files in HDFS, then stores the contents as documents in a MarkLogic Server database. |
| ZipContentLoader | HDFS | MarkLogic Server | Reads text files from zip files in HDFS, then stores the contents as documents in a MarkLogic Server database. |

7.4.1 HelloWorld

This example extracts the first word from all the XML documents in a MarkLogic Server database containing text nodes, sorts the words, concatenates them into a single string, and saves the result as a text document in MarkLogic Server. The example uses basic input mode with the default document selector and subexpression expression. The example uses MarkLogic Server for both input and output.

For detailed instructions on configuring and running this sample, see “Running the HelloWorld Sample Application” on page 19.

Though you can use the sample with any input documents, it is intended to be used with a small data set. It is not optimized for efficient resource use across large data sets. Only XML documents with text nodes contribute to the final results.

This example uses the configuration file:

```
marklogic-hello-world.xml
```

Use the following command to run the example job, with suitable substitution for

```
$CONNECTOR_HOME:
```

```
hadoop jar $CONNECTOR_HOME/lib/marklogic-mapreduce-examples-1.0.jar \
  com.marklogic.mapreduce.examples.HelloWorld
```

To view the results, use Query Console to explore the output database. The sample creates `HelloWorld.txt`. If you use the input data from “Configuring the Job” on page 21, `HelloWorld.txt` should contain the phrase “hello world”.

7.4.2 LinkCountInDoc

This example calculates reference counts for each document in a set of Wikipedia-based documents, and then stores the reference count for each document as a new `<ref-count>` child node of the document. The example uses MarkLogic Server for input and output.

Note: Before running the sample, follow the instructions in “Additional Sample Data Setup” on page 90.

This example uses the configuration file:

```
marklogic-nodein-nodeout.xml
```

Use the following command to run the example job, with a suitable substitution for `$CONNECTOR_HOME`:

```
hadoop jar $CONNECTOR_HOME/lib/marklogic-mapreduce-examples-1.0.jar \
  com.marklogic.mapreduce.examples.LinkCountInDoc \
  marklogic-nodein-nodeout.xml
```

The intra-collection reference counts are stored as new `<ref-count>` elements under the root of each document. Run the following query in Query Console against your XDBC App Server to see a list of documents `ref-count` elements:

```
xquery version "1.0-ml";

for $ref in //ref-count
return fn:concat(xdmp:node-uri($ref), " ", $ref/text())
```

You should see results similar to the following:

```
enwiki/Ayn Rand 1
enwiki/List of characters in Atlas Shrugged 4
enwiki/Academy Award for Best Art Direction 1
enwiki/Academy Award 2
enwiki/Aristotle 5
```

7.4.3 LinkCountInProperty

This example calculates reference counts for each document in a set of Wikipedia-base documents, and then stores the reference count for each document as a property of the document. The examples uses MarkLogic Server for input and output.

Note: Before running the sample, follow the instructions in “Additional Sample Data Setup” on page 90.

This example uses the configuration file:

```
marklogic-textin-propout.xml
```

Use the following command to run the example job, with a suitable substitution for `$CONNECTOR_HOME`:

```
hadoop jar $CONNECTOR_HOME/lib/marklogic-mapreduce-examples-1.0.jar \
  com.marklogic.mapreduce.examples.LinkCountInProperty \
  marklogic-textin-propout.xml
```

The intra-collection reference counts are stored as new `<ref-count>` property elements of each document. Run the following query in Query Console against your XDBC App Server to see a list of documents with at least 20 references:

```
xquery version "1.0-m1";

for $ref in xdmp:document-properties()//ref-count
return fn:concat(xdmp:node-uri($ref), " ", $ref/text())
```

You should see results similar to the following:

```
enwiki/Ayn Rand 1
enwiki/List of characters in Atlas Shrugged 4
enwiki/Academy Award for Best Art Direction 1
enwiki/Academy Award 2
enwiki/Aristotle 5
```

7.4.4 LinkCountValue

This example calculates reference counts for each document in a set of Wikipedia-base documents, and then stores the reference counts in HDFS. The examples uses MarkLogic Server for input and HDFS for output.

Note: Before running the sample, follow the instructions in “Additional Sample Data Setup” on page 90.

This example uses the configuration file:

```
marklogic-textin-textout.xml
```

Use the following command to run the example job, with suitable substitutions for `$CONNECTOR_HOME` and the `HDFS_OUTPUT_DIR`. The `HDFS_OUTPUT_DIR` must not already exist.

```
hadoop jar $CONNECTOR_HOME/lib/marklogic-mapreduce-examples-1.0.jar \
  com.marklogic.mapreduce.examples.LinkCountValue \
  marklogic-textin-textout.xml HDFS_OUTPUT_DIR
```

To view the results, examine the results in `HDFS_OUTPUT_DIR`, as described in “Interacting with HDFS” on page 92. For example, if you use `/home/you/lcv` for `HDFS_OUTPUT_DIR`:

```
$ hadoop dfs -ls /home/you/lcv
... part-r-00000
$ hadoop dfs -cat /home/you/lcv/part-r-00000 | grep "^Aristotle"
Aristotle      5
```

Each topic title is followed by a reference count. The raw output differs from the results for the `LinkCountInDoc` and `LinkCountInProperty` examples because `LinkCountValue` generates counts for all references, rather than only for documents in the database.

7.4.5 LinkCount

This example calculates reference counts for each document in a set of Wikipedia-base documents, and then stores the reference counts in HDFS. The examples uses MarkLogic Server for input and HDFS for output. This example is the same as the `LinkCountValue` example, but it uses advanced input mode instead of basic input mode.

Note: Before running the sample, follow the instructions in “Additional Sample Data Setup” on page 90.

This example uses the configuration file:

```
marklogic-advanced.xml
```

Use the following command to run the example job, with suitable substitutions for

`$CONNECTOR_HOME` and the `HDFS_OUTPUT_DIR`:

```
hadoop jar $CONNECTOR_HOME/lib/marklogic-mapreduce-examples-1.0.jar \
  com.marklogic.mapreduce.examples.LinkCount \
  marklogic-advanced.xml HDFS_OUTPUT_DIR
```

To view the results, examine the results in `HDFS_OUTPUT_DIR`, as described in “Interacting with HDFS” on page 92. For example, if you use `/home/you/lc` for `HDFS_OUTPUT_DIR`:

```
$ hadoop dfs -ls /home/you/lc
... part-r-00000
$ hadoop dfs -cat /home/you/lcv/part-r-00000 | grep "^Aristotle"
Aristotle      5
```

Each topic title is followed by a reference count. The raw output differs from the results for the `LinkCountInDoc` and `LinkCountInProperty` examples because `LinkCount` generates counts for all references, rather than only for documents in the database.

For details on advanced input mode, see “Advanced Input Mode” on page 50.

7.4.6 LinkCountCooccurrences

This example calculates reference counts for each document in a set of Wikipedia-base documents by using an element attribute lexicon, and then stores the reference counts in HDFS. The examples uses MarkLogic Server for input and HDFS for output.

The sample uses `com.marklogic.mapreduce.functions.ElemAttrValueCooccurrences` (a wrapper around `cts:element-attribute-value-co-occurrences`) to find all `href` attributes which occur along with `title` attributes in side anchor tags. The attribute range indexes created in “Additional Sample Data Setup” on page 90 support this operation. The map input key-value pairs are `(Text, Text)` pairs where the key is the `href` and the value is the title.

Note: Before running the sample, follow the instructions in “Additional Sample Data Setup” on page 90.

This example uses the configuration file:

```
marklogic-lexicon.xml
```

Use the following command to run the example job, with suitable substitutions for `$CONNECTOR_HOME` and the `HDFS_OUTPUT_DIR`:

```
hadoop jar $CONNECTOR_HOME/lib/marklogic-mapreduce-examples-1.0.jar \
  com.marklogic.mapreduce.examples.LinkCountCooccurrences \
  marklogic-lexicon.xml HDFS_OUTPUT_DIR
```

To view the results, examine the results in `HDFS_OUTPUT_DIR`, as described in “Interacting with HDFS” on page 92. For example, if you use `/home/you/lc` for `HDFS_OUTPUT_DIR`:

```
$ hadoop dfs -ls /home/you/lcco
... part-r-00000
$ hadoop dfs -cat /home/you/lcco/part-r-00000 | grep "^Aristotle"
Aristotle      5
```

Each topic title is followed by a reference count. The raw output differs from the results for the `LinkCountInDoc` and `LinkCountInProperty` examples because `LinkCountCooccurrences` generates counts for all references, rather than only for documents in the database.

7.4.7 RevisionGrouper

This sample application demonstrates using `KeyValueOutputFormat` and a custom output query. The sample places each document in a collection of Wikipedia articles into a collection, based on the year the article was last revised. The sample uses MarkLogic Server for input and output. The job has no reduce phase.

Note: Before running the sample, follow the instructions in “Additional Sample Data Setup” on page 90

The map function input key-value pairs that are the revision timestamp nodes matching the XPath expression `fn:collection()//wp:revision/wp:timestamp`, using the expression in `mapreduce.marklogic.input.subdocumentexpr`. These nodes are of the form:

```
<timestamp>2007-09-28T08:07:26Z</timestamp>
```

The map function picks the year (2007) off the timestamp and generates output key-value pairs where the key is the document URI as a string and the value is the year as a string. Each pair is then passed to the output query defined in `mapreduce.marklogic.output.query`, which adds the document named in the key to a collection named after the year.

This example uses the configuration file:

```
marklogic-nodein-qryout.xml
```

Use the following command to run the example job, with a suitable substitution for `$CONNECTOR_HOME`:

```
hadoop jar $CONNECTOR_HOME/lib/marklogic-mapreduce-examples-1.0.jar \  
  com.marklogic.mapreduce.examples.RevisionGrouper \  
  marklogic-nodein-qryout.xml
```

To view the results, use Query Console to explore the `hadoop-samples` database. You should see the documents are now in collections based on the year in which they were revised. Alternatively, run a query similar to the following to see a list of documents in the collection for the year 2009:

```
xquery version "1.0-ml";  
for $d in fn:collection("2009")  
return xdmp:node-uri($d)
```

7.4.8 BinaryReader

This sample application demonstrates using the `mapreduce.marklogic.output.bindsplitrange` configuration property with advanced input mode. The sample extracts the first 1K bytes from each (binary) document in a database and saves the result in HDFS. The sample uses MarkLogic Server for input and HDFS for output. The sample has no reduce phase.

The input query defined in `marlogic-subbinary.xml` uses the `splitstart` and `splitend` external variables provided by the MarkLogic Connector for Hadoop to optimize input query performance. For details on this feature, see “Optimizing Your Input Query” on page 53.

The sample requires a database containing one or more binary documents, and an XDBC App Server. Since the sample assumes all documents in the database are binary documents, you should not use the database and content set up in “Additional Sample Data Setup” on page 90.

Follow these steps to set up the `BinaryReader` sample application:

1. Create a MarkLogic Server database to hold the input data.
2. Create an XDBC App Server and attach it to the database created in Step 1. You may use any root.

3. Edit the configuration file `marklogic-subbinary.xml` to configure the job to use the App Server created in Step 2. For details, see “Modify the Sample Configuration Files” on page 88.
4. Run the sample using the following command, substituting an appropriate value for `HDFS_OUTPUT_DIR`:

```
hadoop jar $CONNECTOR_HOME/lib/marklogic-mapreduce-examples-1.0.jar \
  com.marklogic.mapreduce.examples.BinaryReader \
  marklogic-subbinary.xml HDFS_OUTPUT_DIR
```

If using standalone Hadoop, view the results in `HDFS_OUTPUT_DIR`. If using pseudo-distributed or fully-distributed Hadoop, view the results using the Hadoop `dfs` command. For example:

```
$ hadoop dfs -ls HDFS_OUTPUT_DIR
```

7.4.9 ContentReader

This sample application writes documents in a MarkLogic Server database to the HDFS file system, using an SSL-enabled connection to MarkLogic Server. The input database is the database associated with your XDBC App Server. This sample uses MarkLogic Server for input and HDFS for output.

Note: This sample copies the entire contents of the database to HDFS. Choose a target database accordingly, or modify the sample’s configuration file to limit the selected documents.

Before running the sample, install and modify the sample configuration files and load the sample data, as discussed in “Modify the Sample Configuration Files” on page 88. This example uses the configuration file:

```
marklogic-docin-textout.xml
```

Before running the sample, configure your XDBC App Server and Java environment to use SSL. For details, see “Making a Secure Connection to MarkLogic Server with SSL” on page 13. You might need to import the MarkLogic Server self-signed certificate into your JRE default keystore using the Java `keytool` utility. See the Java documentation for details on adding certificates to the default keystore.

Run the sample as follows, substituting appropriate paths for `$CONNECTOR_HOME` and `HDFS_OUTPUT_DIR`. The output directory must not already exist.

```
hadoop jar $CONNECTOR_HOME/lib/marklogic-mapreduce-examples-1.0.jar \
  com.marklogic.mapreduce.examples.ContentReader \
  marklogic-docin-textout.xml HDFS_OUTPUT_DIR
```

If using standalone Hadoop, view the results in `~/HDFS_OUTPUT_DIR`. If using pseudo-distributed or fully-distributed Hadoop, view the results using the Hadoop `dfs` command. For example:

```
$ hadoop dfs -ls HDFS_OUTPUT_DIR
$ hadoop dfs -tail HDFS_OUTPUT_DIR/part-m-00000
```

7.4.10 ContentLoader

This sample application loads files in an HDFS directory into a MarkLogic Server database as documents. The destination database is the database associated with your XDBC App Server. The sample uses HDFS for input and MarkLogic Server for output. This sample is a map-only job. That is, there is no reduce step.

Before running the sample, install and modify the sample configuration files, as discussed in “Modify the Sample Configuration Files” on page 88. This example uses the configuration file:

```
marklogic-textin-docout.xml
```

Follow these steps to run the `ContentLoader` sample application:

1. Select the text, XML, or binary files to be loaded into the database.
2. If using Hadoop standalone, create an input directory in your home directory to hold the input files. For example:


```
$ mkdir ~/input_dir
```
3. If using Hadoop pseudo- or fully-distributed, create an input directory in HDFS to hold the input files. For example:


```
$ hadoop dfs -mkdir input_dir
```
4. Copy the input files into the input directory created in Step 2 or Step 3. For example:


```
$ cp your_input_files ~/input_dir # standalone
$ hadoop dfs -put your_input_files input_dir # distributed
```
5. If your input content is not XML, edit `$HADOOP_HOME/conf/marklogic-textin-docout.xml` to set the output content type. For example, to load binary files, add:

```
<property>
  <name>mapreduce.marklogic.output.content.type</name>
  <value>binary</value>
</property>
```

6. Run the sample application, substituting appropriate paths for `$CONNECTOR_HOME` and `input_dir`:

```
hadoop jar $CONNECTOR_HOME/lib/marklogic-mapreduce-examples-1.0.jar \
  com.marklogic.mapreduce.examples.ContentLoader \
  marklogic-textin-docout.xml input_dir
```

- Using Query Console, explore the database associated with your XDBC App Server and observe that the input files appear as documents in the database.

The document URIs in the database correspond to the HDFS path. For example, if one of the input documents is located in HDFS on `samples.marklogic.com` as `/user/guest/data/file1.xml`, then the document URI in the database is:

```
hdfs://samples.marklogic.com/user/guest/data/file1.xml
```

If you receive an error similar to the following, then you must change the directory creation database configuration setting to “manual”.

```
java.lang.IllegalStateException: Manual directory creation mode is
required.
```

7.4.11 ZipContentLoader

This sample application loads the contents of zip files in an HDFS directory into a MarkLogic Server database as documents. The destination database is the database associated with your XDBC App Server. The sample uses HDFS for input and MarkLogic Server for output. This sample is a map-only job. That is, there is no reduce step.

Before running the sample, install and modify the sample configuration files, as discussed in “Modify the Sample Configuration Files” on page 88. This example uses the configuration file:

```
marklogic-textin-docout.xml
```

Follow these steps to run the `zipContentLoader` sample application:

- Create one or more zip files containing text, XML, or binary files.
- If using Hadoop standalone, create an input directory in your home directory to hold the zip input files. For example:


```
$ mkdir ~/zip_input_dir
```
- If using Hadoop pseudo- or fully-distributed, create an input directory in HDFS to hold the input files. For example:


```
$ hadoop dfs -mkdir zip_input_dir
```
- Copy the input zip files into the input directory created in Step 2 or Step 3. For example:


```
$ cp your_input_files ~/zip_input_dir # standalone
$ hadoop dfs -put your_data.zip zip_input_dir # distributed
```
- If your zip file content is not XML, set the output content type in `$HADOOP_HOME/conf/marklogic-textin-docout.xml`. For example, to load binary files, add:

```
<property>
  <name>mapreduce.marklogic.output.content.type</name>
  <value>binary</value>
</property>
```

6. Run the sample application, substituting appropriate paths for `$CONNECTOR_HOME` and `zip_input_dir`:

```
hadoop jar $CONNECTOR_HOME/lib/marklogic-mapreduce-examples-1.0.jar \
  com.marklogic.mapreduce.examples.ZipContentLoader \
  marklogic-textin-docout.xml zip_input_dir
```

7. Using Query Console, explore the database associated with your XDBC App Server and observe that the zip file contents appear as documents in the database.

The document URIs in the database correspond to the paths within the zip file. For example, if the zip file contents are rooted at a folder named “enwiki”, and that folder contains a file named “Wisconsin”, then the resulting document URI is:

```
enwiki/Wisconsin
```

If you receive an error similar to the following, use the Admin Interface to change the directory creation database configuration setting to “manual”:

```
java.lang.IllegalStateException: Manual directory creation mode is
required.
```

8.0 Technical Support

MarkLogic provides technical support according to the terms detailed in your Software License Agreement or End User License Agreement. For evaluation licenses, MarkLogic may provide support on an “as possible” basis.

For customers with a support contract, we invite you to visit our support website at <http://support.marklogic.com> to access information on known and fixed issues.

For complete product documentation, the latest product release downloads, and other useful information for developers, visit our developer site at <http://developer.marklogic.com>.

If you have questions or comments, you may contact MarkLogic Technical Support at the following email address:

support@marklogic.com

If reporting a query evaluation problem, please be sure to include the sample XQuery code.

Product Notices

COPYRIGHT

Copyright © 2012 MarkLogic Corporation. All rights reserved.

The MarkLogic software is protected by United States and international copyright laws, and incorporates certain third party libraries and components which are subject to the attributions, terms, conditions and disclaimers set forth below.

1. Contains SAP BusinessObjects Text Analysis XI from SAP AG. Copyright © 1996-2011. All rights reserved.
2. Highslide Software from Highslide Software Torstein Honsi. All Highslide Software is protected by local and international copyright laws. All rights reserved.
3. Icons developed by Yusuke Kamiyamane. Copyright © 2011 Yusuke Kamiyamane. All rights reserved. Icons are licensed subject to <http://creativecommons.org/licenses/by/3.0/legalcode>.
4. Antenna House OfficeHTML Copyright © 2000-2008 Antenna House, Inc. All rights reserved.
5. Argus Copyright ©1999-2008 Icenit Technology Ltd. All rights reserved.
6. Rosette Linguistics Platform 6.5.2/6.5.3 from Basis Technology Corporation, Copyright © 2004-2008 Basis Technology Corporation. All rights reserved.
7. ISYS Search. Copyright © 2011 ISYS™ Search Software, Inc. All rights reserved.
8. Software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>) Copyright © 1995-1998 Eric Young (eay@cryptsoft.com). All rights reserved. Copyright © 1998-2011 The OpenSSL Project. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. 3. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)" 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org. 5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project. 6. Redistributions of any form whatsoever must retain the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)" THIS SOFTWARE

IS PROVIDED BY THE OpenSSL PROJECT ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com). Additional terms may apply to the foregoing software as further set forth at <http://www.openssl.org/source/license.html>

9. ICU v. 4.2.1 available from <http://site.icu-project.org/> Copyright © 1995-2011 International Business Machines Corporation and others. All rights reserved. Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation. THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE. Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

10. Tidy available at <http://www.w3.org/People/Raggett/tidy/>. Copyright © 1998-2008 World Wide Web Consortium (Massachusetts Institute of Technology, European Research Consortium for Informatics and Mathematics, Keio University). All Rights Reserved.

11. 'zlib' general purpose compression library v. 1.2.3. Copyright © 1995-2010 Jean-loup Gailly and Mark Adler. http://www.zlib.net/zlib_license.html

12. RSA Data Security, Inc. MD5 Message-Digest Algorithm, copyright © 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved. License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing this software or this function. License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work. RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind. These notices must be retained in any copies of any part of this documentation and/or software.

13. TRE software, available at <http://laurikari.net/tre/>. Copyright © 2001-2009 Ville Laurikari <vl@iki.fi>. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: (1) Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. (2) Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

14. Eigen Library, version 2.0.10, available at http://eigen.tuxfamily.org/index.php?title=Main_Page. Eigen Library is free software: you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation, either version 3 of the License, or (at your option) any later version. Eigen Library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details. You should have received a copy of the GNU Lesser General Public License along with Eigen Library. If not, see: <http://www.gnu.org/copyleft/lesser.html>

15. FunctX XQuery Library v.1.0 available at <http://www.xqueryfunctions.com/xq/>, Copyright © Datypic (Priscilla Walmsley). You can redistribute such library and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version. This library is distributed in the

hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License as available at <http://www.gnu.org/licenses/lgpl-2.1.html#SEC4> for more details.

16. DeployJava.js v. 1.8, Copyright © 2006, 2011, Oracle and/or its affiliates. All rights reserved. ORACLE PROPRIETARY/CONFIDENTIAL. Use of the foregoing library is subject to the following license terms. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. Neither the name of Oracle nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

17. YUI libraries v.2.8 available at <http://developer.yahoo.com/yui/>. Copyright © 2011 Yahoo! Inc. All rights reserved. Use of the foregoing libraries is subject to the following license terms. Redistribution and use of this software in source and binary forms, with or without modification, are permitted provided that the following conditions are met: Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. Neither the name of Yahoo! Inc. nor the names of YUI's contributors may be used to endorse or promote products derived from this software without specific prior written permission of Yahoo! Inc. THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING

NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. Additional notices related to components included in the foregoing library are also available at the following URL: <http://yuilibrary.com/license/>

18. XSLTForms libraries available at <http://sourceforge.net/projects/xsltforms/>. You can redistribute such libraries and/or modify them under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version. These libraries are distributed in the hope that they will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License as available at <http://www.gnu.org/licenses/lgpl-2.1.html#SEC4> for more details.

19. jQuery v.1.7 libraries available at <http://jquery.com/> and jQuery UI v.1.8.2 libraries available at <http://jqueryui.com/>. Copyright © 2011 John Resig. Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software. THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

20. JSON v. 2 libraries available at <http://www.json.org/json2.js>. Copyright © 2002 JSON.org. Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software. The Software shall be used for Good, not Evil. THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

21. CodeMirror libraries available at <http://codemirror.net/>. Copyright © 2011 by Marijn Haverbeke marijnh@gmail.com. Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software. **THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.**

22. cURL v. 7.22.0 available at <http://curl.haxx.se/>. Copyright © 1996 - 2011, Daniel Stenberg, <daniel@haxx.se>. All rights reserved. Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies. **THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.** Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

23. History.js available <https://github.com/balupton/History.js/> Copyright © 2011, Benjamin Arthur Lupton. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. Neither the name of Benjamin Arthur Lupton nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. **THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,**

PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

24. Sample content provided for use with the software is copyrighted by the respective authors as contributed by each of them to Wikipedia pursuant to the Creative Commons Attribution-ShareAlike License found at <http://creativecommons.org/licenses/by-sa/3.0/>.

25. Visual Studio Tools for Office 3.0 (Microsoft VSTO 3.0) licensed under Microsoft Public License Microsoft Public License (Ms-PL) found at <http://msdn.microsoft.com/en-us/library/ff647676.aspx>.

26. The MarkLogic Connector for Hadoop contains Apache Jakarta Commons Modeler. Copyright 2001-2007 The Apache Software Foundation. Licensed under the Apache License, Version 2.0 (the "License"). You may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0>. Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

27. jemalloc v 2.2.4. <http://www.canonware.com/jemalloc/index.html>, released under the terms of the following BSD-derived licenses: Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the above copyright notice(s), this list of conditions and the following disclaimer. 2. Redistributions in binary form must reproduce the above copyright notice(s), this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER(S) ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER(S) BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. Portions of the software are copyright (C) 2002-2011 Jason Evans, and portions copyright (C) 2009-2011 Facebook, Inc. and copyright (C) 2007-2010 Mozilla Foundation. Neither the name of Facebook, Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. All rights reserved.

28. GzipHelper; version: 1.0, Feb, 2003, authored by Gao Dasheng, is distributed subject to the following terms: This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software. Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions: (1.) the origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required; (2) altered source versions must be plainly marked as such, and must not be misrepresented as being the original software; (3) this notice may not be removed or altered from any source distribution.

TRADEMARK NOTICE

The MarkLogic name and logo are registered trademarks of MarkLogic Corporation in the United States and/or other countries. Excel and PowerPoint are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Any other names or logos included in the Software, this notice or the documentation are property of the respective trademark owners.