

HITACHI
Inspire the Next



Pentaho and Big Data Ingestion Patterns

HITACHI

Inspire the Next

Change log (if you want to use it):

Date	Version	Author	Changes

Contents

- Overview 1
 - Before You Begin..... 1
 - Use Case: Many Small Text Files 1
- Data Ingestion to Big Data 2
 - Data from External JDBC Databases..... 3
 - Using Sqoop to Ingest Data 3
 - Using Pentaho JDBC to HDFS 4
 - Using PMR JDBC to HDFS 4
 - Data from Local/External File Sources 5
 - Using Pentaho File to HDFS 5
 - Use PMR File to HDFS 6
 - Data from Realtime Systems 7
- Data Processing in Big Data 8
 - Text-Based Data in Large Files 9
 - Processing Large Files from HDFS with PMR..... 9
 - Processing Large Files from HDFS with AEL Spark 9
 - Text-Based Data in Many Small Files 10
- Related Information 11

This page intentionally left blank.

Overview

This document covers some best practices on collecting Pentaho Data Integration (PDI) and Hadoop usage patterns for data ingestion and processing.

Our intended audience is Pentaho administrators, or anyone with a background in Hadoop and data ingestion who is interested in choosing the best methods for processing data.

The intention of this document is to speak about topics generally; however, these are the specific versions covered here:

The [Components Reference](#) in Pentaho Documentation has a complete list of supported software and hardware.

Before You Begin

Before beginning, use the following information to prepare for the procedures described in the main section of the document.

This document assumes that you have knowledge of Pentaho and Hadoop and already know what format your data will arrive in.

Use Case: Many Small Text Files

Janice has many small files of product descriptions in Excel spreadsheets. She needs to load this data as quickly as possible so that Marketing can redo the product catalog.

Janice decides to use a data loading pattern of a distributed load process, with the files processed in parallel by different Hadoop nodes.

Data Ingestion to Big Data

Data ingestion is the process of getting data from external sources into big data. There are different patterns that can be used to load data to Hadoop using PDI. The preferred ingestion format for landing data from Hadoop is Avro.

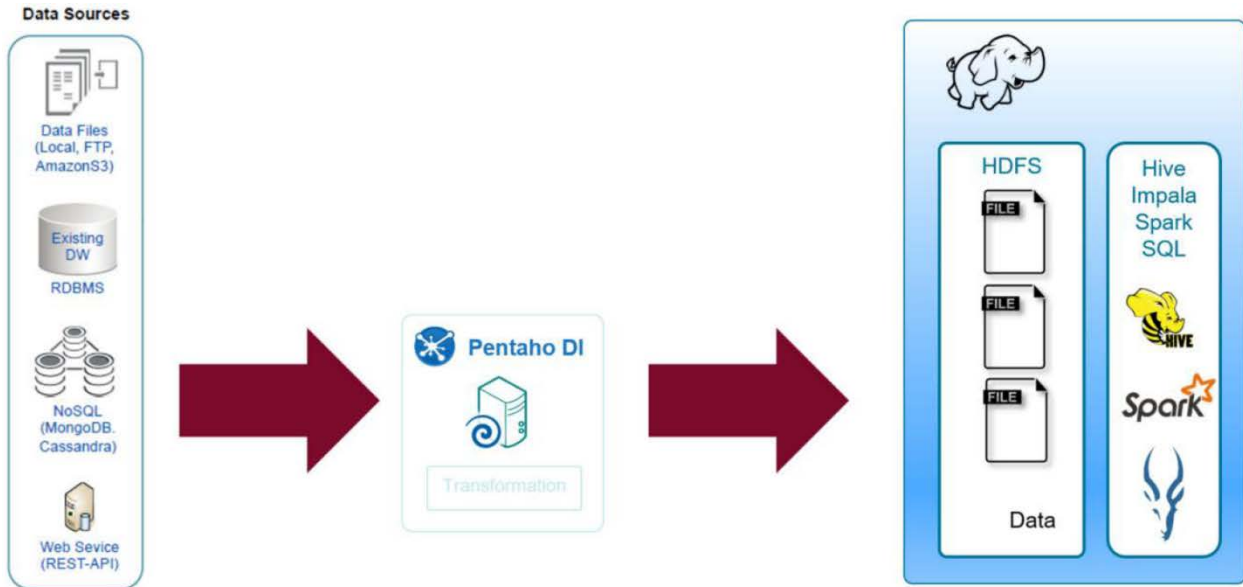


Figure 1: Ingestion to Big Data

You can find details on these topics in the following sections, sorted by where your data is arriving from:

- [Data from External JDBC Databases](#)
- [Data from Local/External File Sources](#)
- [Data from Realtime Systems](#)

Data from External JDBC Databases

When you collect source data in external Java Database Connectivity (JDBC) databases, you have multiple methods you can choose from to ingest data:

Table 1: Ingestion from External JDBC Databases

Situation	Available Method(s)
Data is accessible from any Hadoop node. Nodes have connectivity to the external database. No transformation/change is needed.	Use Sqoop to ingest data.
Data is accessible from any Hadoop node. Nodes have connectivity to the external database. Data may or may not need transformation.	Use Pentaho MapReduce (PMR) JDBC to Hadoop Distributed File System (HDFS) OR Use PDI Cluster on YARN. (See Pentaho Big Data On-Cluster Processing.)
Data is not accessible from Hadoop nodes. Data is only accessible by edge nodes or Pentaho Server.	Use Pentaho JDBC to HDFS OR Use Pentaho distributed processing (PDI Cluster).

Using Sqoop to Ingest Data

Using Sqoop to ingest your data involves a distributed load process, with multiple connections to the source database from different Hadoop nodes.

The HDFS destination for your text can be text, Avro, or Parquet. Avro schemas are created automatically by Sqoop based on source data. You can also enable different Avro or Parquet options if you wish.

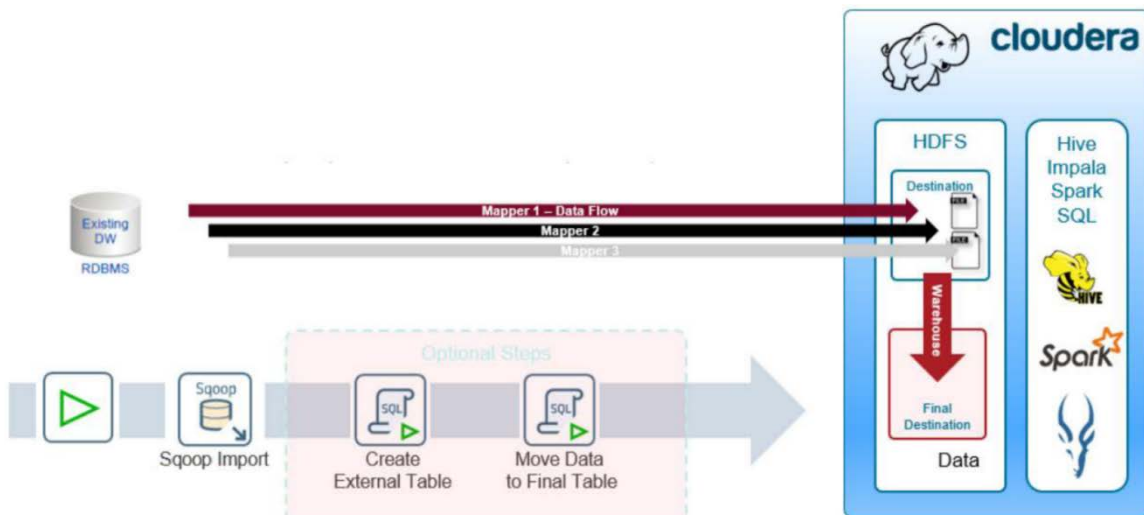


Figure 2: Using Sqoop to Ingest Data

Using Pentaho JDBC to HDFS

In this case, you can have a single worker process the entire dataflow using PDI (Kitchen/Pan/Carte). The HDFS destination for your text can be text, Parquet, or Avro¹. You can also enable different Avro or Parquet options if you wish.

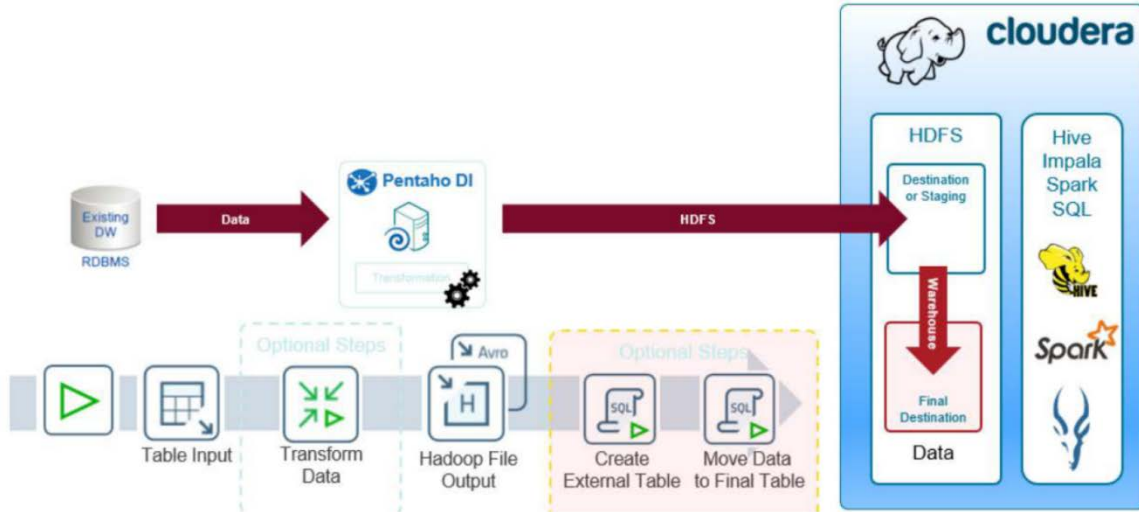


Figure 3: Pentaho JDBC to HDFS

Using PMR JDBC to HDFS

This method uses a distributed load process, with multiple connections to source databases from different Hadoop nodes. You must define how to segment the data ranges. The HDFS destination for your text can be text, Parquet, or Avro. You can also enable different Parquet options if you wish.

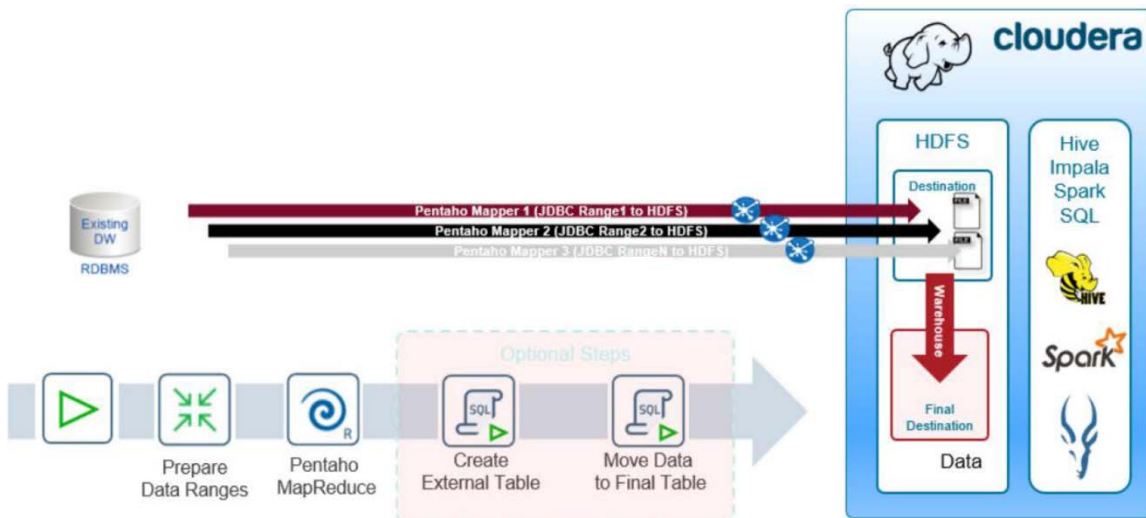


Figure 4: PMR JDBC to HDFS

Non-nested (plain) Parquet and Avro are supported.¹

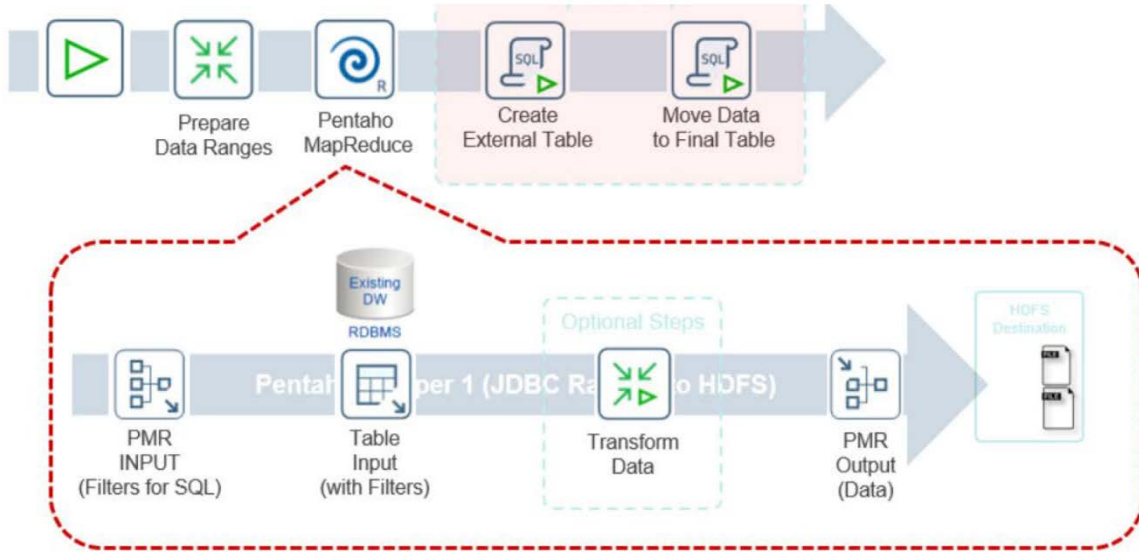


Figure 5: PMR Breakdown with Optional Steps

Data from Local/External File Sources

When you store source data in local/external filesystems, you may find one of these patterns useful:

Table 2: Ingestion from Local/External File Sources

Situation	Available Method(s)
Data is not accessible from Hadoop nodes. Data is only accessible by edge nodes or Pentaho Server	Use Pentaho File to HDFS.
Data is accessible from any Hadoop node. Nodes have connectivity to the external database.	Use PMR File to HDFS.

Using Pentaho File to HDFS

In this case, you can have a single worker process the entire dataflow using PDI (Kitchen/Pan/Carte). The files are processed sequentially.

The HDFS destination for your text can be text, Parquet, or Avro². You can also enable different Avro or Parquet options if you wish.

²Non-nested (plain) Parquet and Avro are supported.

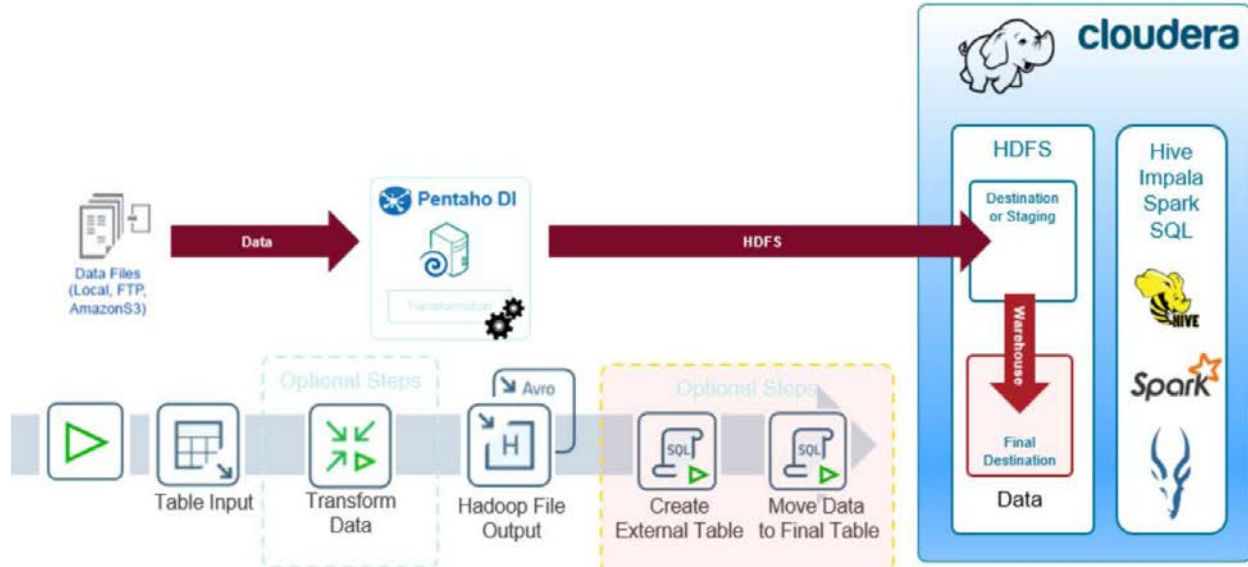


Figure 6: Pentaho File to HDFS

Use PMR File to HDFS

This method uses a distributed load process to distribute the file loading in parallel. The set of files is processed by different Hadoop nodes.

The HDFS destination for your text can be text, Parquet, or Avro. You can also enable different Avro or Parquet options if you wish.

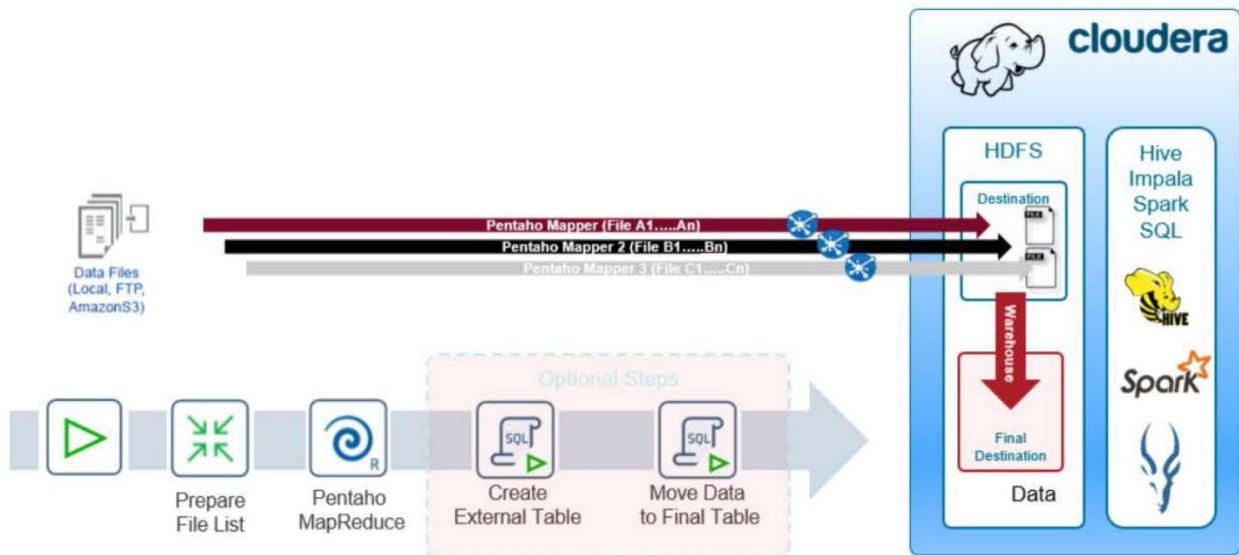


Figure 7: PMR File to HDFS

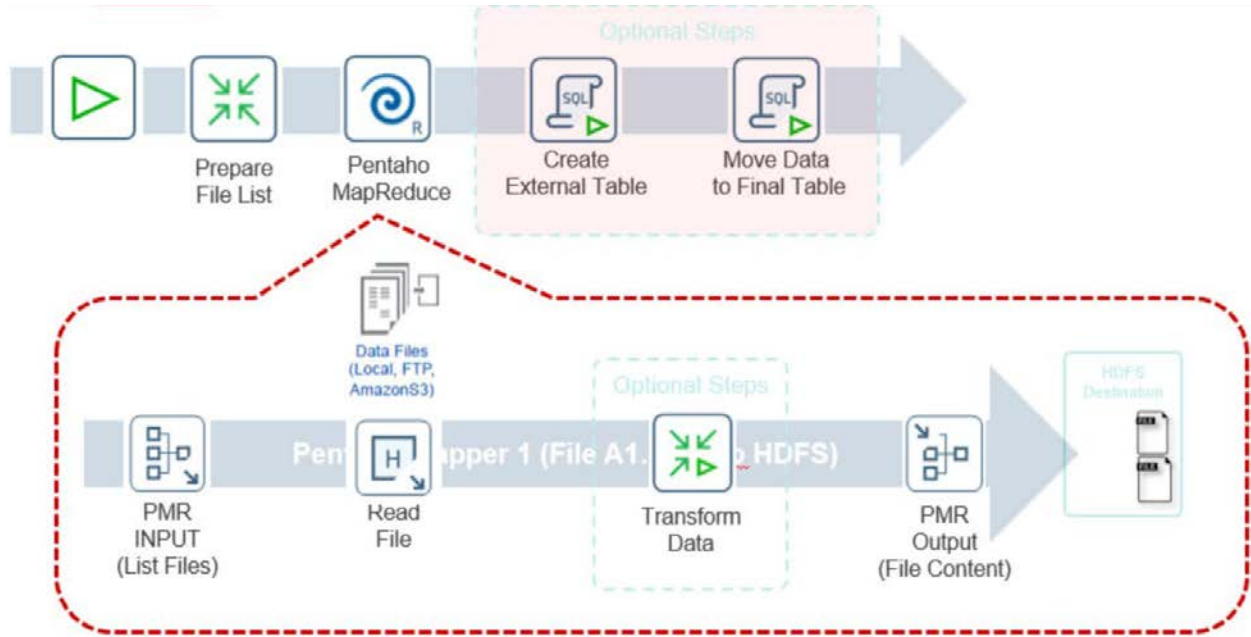


Figure 8: PMR Breakdown with Optional Steps

Data from Realtime Systems

In cases where your data comes from realtime systems like messages or streaming sources, you have several pattern options available to you. See [Realtime Data Processing with PDI](#) for more information.

Data Processing in Big Data

Several patterns are available for processing data stored in Hadoop that is accessible with HDFS, Hive, Impala, and so forth.

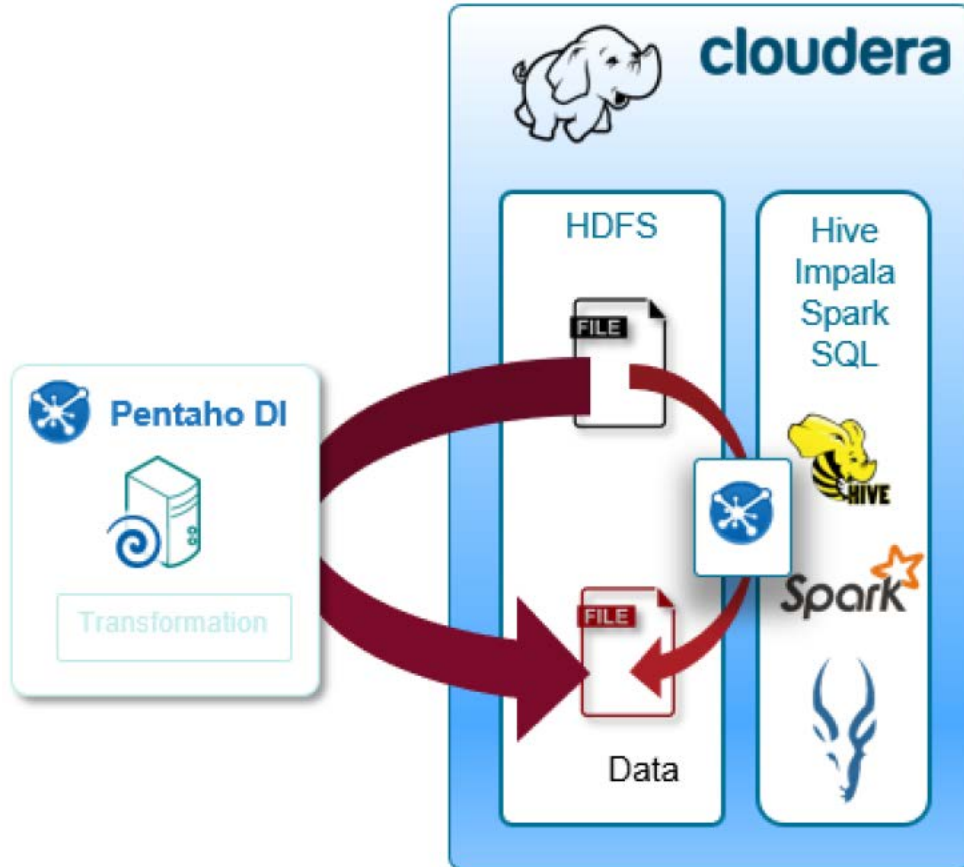


Figure 9: Data Processing

Input and output format may vary from case to case, and different patterns apply best to each case. For regular data processing, you may want to clean, transform, and/or extend your data. For staging your data, you may want to do things like create external Hive definitions on top of your staging data, or move data from staging to a final destination, with proper partitioning and file formatting.

When you need to transform, clean, classify, extend, or filter data that is stored in HDFS, you need to maximize data locality usage.

You can find details on these topics in the following sections:

- [Text-Based Data in Large Files](#)
- [Text-Based Data in Many Small Files](#)

Text-Based Data in Large Files

When your text-based input is stored in large files like XML, JSON, or text, patterns that may work for you are to process with PMR or with Adaptive Execution Layer (AEL) Spark.

Processing Large Files from HDFS with PMR

This processing pattern uses data locality. The source can be text, XML, JSON, or Avro. The HDFS destination for your data can be text, Parquet, or Avro. You can also enable different Avro or Parquet options if you wish.

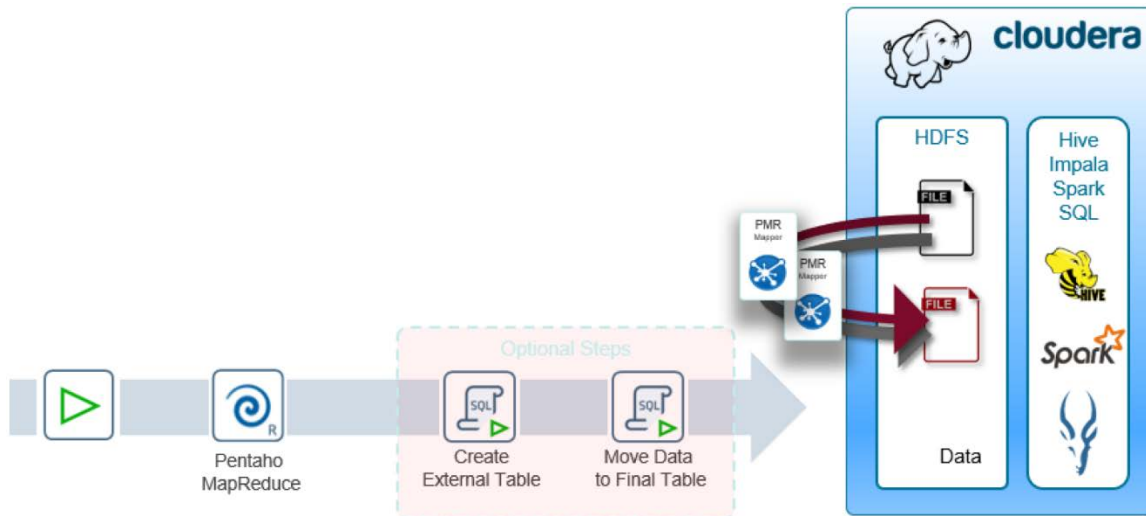


Figure 10: Processing Large Files with PMR

Processing Large Files from HDFS with AEL Spark

This processing pattern uses data locality. The source can be text, XML, or JSON. The HDFS destination for your data will be text.



Figure 11: Processing Large Files with AEL Spark

Text-Based Data in Many Small Files

This pattern uses a distributed load process to distribute the file loading in parallel. The set of files is processed by different Hadoop nodes. This method does not use data locality. See [Parsing XML on PDI](#) for more information.

The data source can be text, XLS/XLSX, or Avro, with other formats supported by Pentaho input. The destination can be text, Parquet, or Avro. You can also enable different Avro or Parquet options, if you wish.

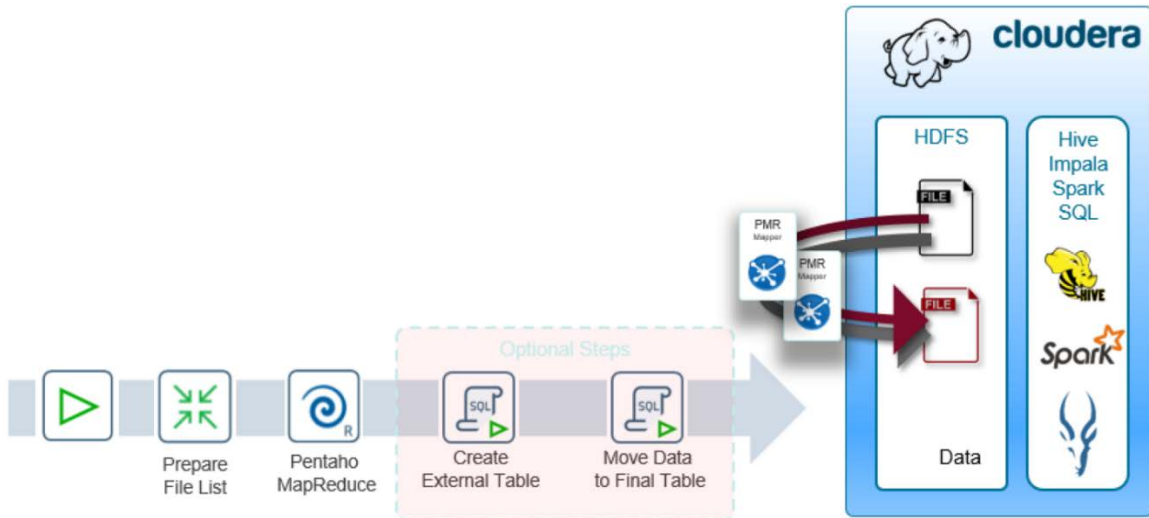


Figure 12: Processing Small Files Using Distributed Load Process

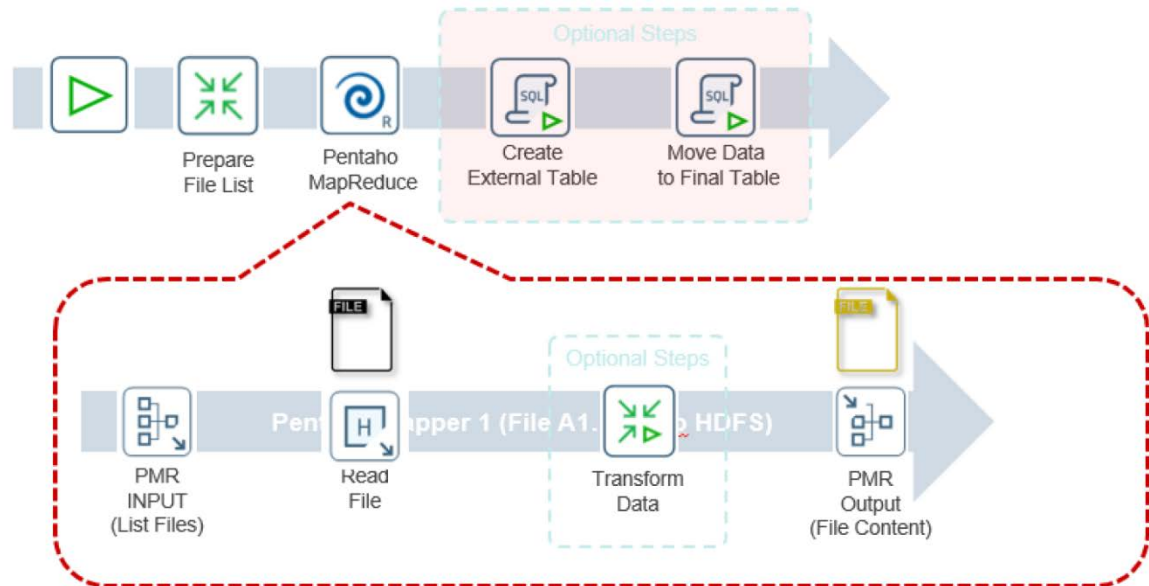


Figure 13: PMR Breakdown with Optional Steps

Related Information

Here are some links to information that you may find helpful while using this best practices document:

- [Big Data and Pentaho](#)
- [Components Reference](#)
- [Realtime Data Processing with PDI](#)
- [Set Up Pentaho to Connect to an Amazon EMR Cluster](#)