



Chapter 7:

PIG, HIVE, HBASE and ZOOKEEPER

1



Lebanese University
Faculty of Information 1

WHAT IS PIG?

Framework for analyzing large un-structured and semi-structured data on top of hadoop

Pig engine: runtime environment where the program executed.

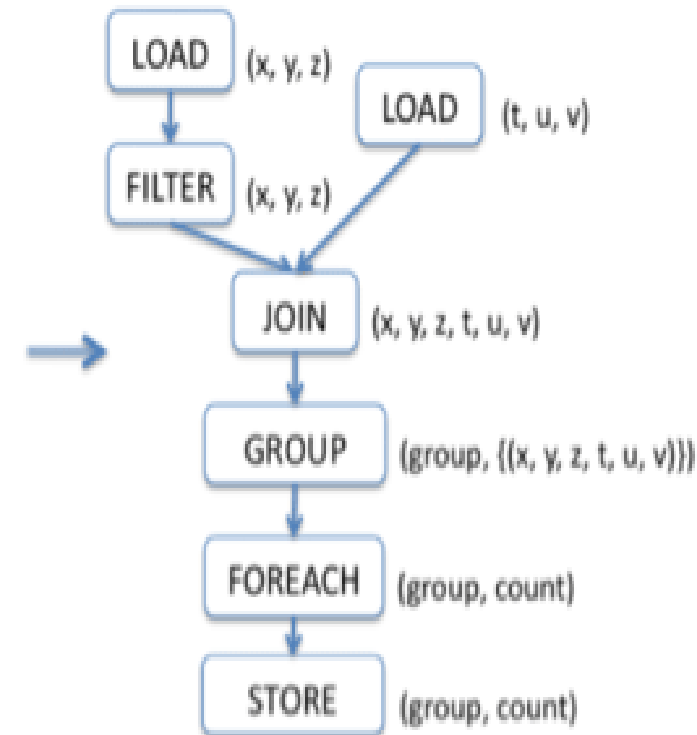
Pig Latin: is simple but powerful data flow language similar to scripting language.

1. SQL
2. Provide common data operations(load, filters, joins, group, store)

Pig Latin

```
A = LOAD 'file1' AS (x, y, z);  
B = LOAD 'file2' AS (t, u, v);  
C = FILTER A by y > 0;  
D = JOIN C BY x, B BY u;  
E = GROUP D BY z;  
F = FOREACH E GENERATE  
    group, COUNT(D);  
STORE F INTO 'output';
```

Logical Plan



PIG CHARACTERISTIC

- A platform analyzing large data sets that runs on top hadoop
- Provides a high-level language for expressing data analysis
- Uses both Hadoop Distributed File System (HDFS) read and write files and MapReducer (execute jobs)

EXECUTION TYPES

- Local Mode:
 - Pig runs in a single JVM and accesses the local filesystem. This mode is suitable only for small datasets and when trying out Pig.
- MapReduce Mode:
 - Pig translates queries into MapReduce jobs and runs them on a Hadoop cluster. It is what you use when you want to run Pig on large datasets.


RUNNING PIG PROGRAMS

There are three ways of executing Pig programs

- Script -> Pig can run a script file that contains Pig commands
- Grunt -> Interactive shell for running Pig commands
- Embedded -> users can run Pig programs from Java using the PigServer class

PIG LATIN DATA FLOW

- A LOAD statement to read data from the system
- A series of “transformation” statement to process the data
- A DUMP statement to view results or STORE statement to save the result



WHAT IS HIVE?

- **Hive:** A data warehousing system to store structured data on Hadoop file system
- Developed by Facebook
- Provides an essay query by executing Hadoop MapReduce plans
- Provides SQL type language for querying called HiveQL or HQL
- The Hive shell is the primary way that we will interact with Hive

HIVE DATA MODEL

- **Tables:** all the data is stored in a directory in HDFS
 - Primitives: numeric, boolean, string and timestamps
 - Complex: Arrays, maps and structs
- **Partitions:** divides a table into parts
 - Queries that are restricted to a particular date or set of dates can run much more efficiently because they only need to scan the files in the partitions that the query pertains to
- **Buckets:** data in each partitions is divided into buckets
 1. Enable more efficient queries
 2. Makes sampling more efficient

MAJOR COMPONENTS OF HIVE

- **UI:** Users submits queries and other operations to the system
- **Driver:** Session handles, executes and fetches APIs modeled on JDBC/ODBC interfaces
- **Metastore:** Stores all the structure information of the various tables and partitions in the warehouse
- **Compiler:** Converts the HiveQL into a plan for execution
- **Execution Engine:** Manages dependencies between these different stages of the plan and executes these stages on the appropriate system components

HIVE VS PIG

PIG

- Procedural Data Flow Language
- For Programming
- Mainly used by Researchers & Programmers
- Operates on the client side of a cluster
- Better dev environments, debuggers expected

HIVE

- Declarative SQL Language
- For creating reports
- Mainly used by Data/Business Analysts
- Operates on the server side of a cluster
- Better integration with technologies expected

HBase

- HBase is a distributed column-oriented database built on top of the HDFS. It is an open-source project and horizontally scalable.
- HBase is a data model that is similar to Google's big table that designed to provide quick random access to huge amounts of structured data.
- HBase is a part of Hadoop ecosystem that provides real-time read/write access to data in the Hadoop File System.
- **HBase stores its data in HDFS.**

Storage Mechanism in HBase

HBase is a **column-oriented database** and the tables in it are sorted by row. The table schema defines only column families, which are the key value pairs. A table have multiple column families and each column family can have any number of columns. Subsequent column values are stored contiguously on the disk. Each cell value of the table has a timestamp. In short, in an HBase:

- Table is a collection of rows.
- Row is a collection of column families.
- Column family is a collection of columns.
- Column is a collection of key value pairs.

Hbase is Not

- Not an SQL Database.
- Not relational.
- No joins.
- No fancy query language and no sophisticated query engine.

Features of HBase

HBase is linearly scalable.

It has automatic failure support.

It provides consistent read and writes.

It integrates with Hadoop, both as a source and a destination.

It has easy java API for client.

It provides data replication across clusters.

When Should I Use HBase?

make sure you have enough data. If you have hundreds of millions or billions of rows, then HBase is a good candidate.

What is Zookeeper?

- ZooKeeper is an open source Apache™ project that provides a centralized infrastructure and services that enable synchronization across a cluster. ZooKeeper maintains common objects needed in large cluster environments. Examples of these objects include configuration information, hierarchical naming space, and so on. Applications can leverage these services to coordinate distributed processing across large clusters.
- Created by Yahoo! Research
- Started as sub-project of Hadoop

What is Zookeeper meant for?

- **Naming service** – Identifying the nodes in a cluster by name. It is similar to DNS, but for nodes
- **Configuration management** – Latest and up-to-date configuration information of the system for a joining node
- **Cluster management** – Joining / leaving of a node in a cluster and node status at real time
- **Leader election** – Electing a node as leader for coordination purpose
- **Locking and synchronization service** – Locking the data while modifying it. This mechanism helps you in automatic fail recovery while connecting other distributed applications like Apache HBase

How Zookeeper Works

- Zookeeper gives you a set of tools to build distributed applications that can safely handle partial failures
- Zookeeper is simple
- Zookeeper is expressive
- Zookeeper facilitates loosely coupled interactions
- Zookeeper is a library
- Zookeeper requires Java to run
- Maintains configuration information