



HBase Architecture & Data Model

Core Components

HRegionServer	Hosts and manages HRegions; handles read/write requests.
HMaster	Assigns regions to RegionServers, handles schema changes, and performs administrative tasks.
ZooKeeper	Maintains configuration information, naming, and distributed synchronization.
HDFS	Hadoop Distributed File System; stores HBase data persistently.

Data Model

HBase is a distributed, scalable, big data store.

It's designed to store and retrieve data within large tables. Data is stored as key-value pairs.

- Table:** Contains rows.
- Row Key:** Unique identifier for a row.
- Column Family:** Groups related columns.
- Column Qualifier:** Identifies a column within a column family.
- Version:** Each cell value has a version timestamp.

Key Concepts

Regions	Tables are split into regions. A region contains a subset of rows. Regions are the unit of distribution and scalability.
Store	Each region contains one or more stores. A store contains a MemStore and zero or more StoreFiles (HFiles).
MemStore	In-memory buffer that stores recent writes.
HFile	Sorted key-value pairs stored on HDFS.

Basic HBase Operations (CLI)

Table Management

<code>create 'table_name', 'column_family1', 'column_family2'</code>	- Creates a new table.
<code>list</code>	- Lists all tables.
<code>describe 'table_name'</code>	- Describes the table schema.
<code>disable 'table_name'</code>	- Disables a table.
<code>enable 'table_name'</code>	- Enables a table.
<code>drop 'table_name'</code>	- Drops a table (must be disabled first).
<code>exists 'table_name'</code>	- Checks if table exists.

Data Manipulation

<code>put 'table_name', 'row_key', 'column_family:qualifier', 'value'</code>	- Inserts or updates a cell value.
<code>get 'table_name', 'row_key'</code>	- Retrieves all columns for a given row.
<code>get 'table_name', 'row_key', 'column_family:qualifier'</code>	- Retrieves a specific cell.
<code>scan 'table_name'</code>	- Scans the entire table.
<code>scan 'table_name', {STARTROW => 'row_key1', STOPROW => 'row_key2'}</code>	- Scans a range of rows.
<code>delete 'table_name', 'row_key', 'column_family:qualifier'</code>	- Deletes a specific cell.
<code>deleteall 'table_name', 'row_key'</code>	- Deletes all cells in a row.

HBase Shell Commands

Namespace Management

<code>create_namespace 'namespace_name'</code>	- Creates a namespace.
<code>list_namespace</code>	- Lists all namespaces.
<code>describe_namespace 'namespace_name'</code>	- Describes a namespace.
<code>alter_namespace 'namespace_name', {METHOD => 'set', 'PROPERTY_NAME' => 'property_value'}</code>	- Alters a namespace property.
<code>drop_namespace 'namespace_name'</code>	- Drops a namespace.

Advanced Scan Operations

<code>scan 'table_name', {COLUMNS => ['col_family1', 'col_family2:qualifier']}</code>	- Scans specific column families or columns.
<code>scan 'table_name', {LIMIT => 10}</code>	- Limits the number of rows returned.
<code>scan 'table_name', {REVERSED => true}</code>	- Scans in reverse order.
<code>scan 'table_name', {FILTER => "RowFilter(=, 'binary:row_key')"</code>	- Applies a filter to the scan.
<code>scan 'table_name', {VERSIONS => 5}</code>	- Retrieves the last 5 versions of each cell.
<code>count 'table_name'</code>	- Counts the number of rows in a table.

Configuration

HBase configuration is managed through <code>hbase-site.xml</code> . Key properties include:
<ul style="list-style-type: none"> <code>hbase.rootdir</code>: HDFS directory for HBase data. <code>hbase.zookeeper.quorum</code>: List of ZooKeeper servers. <code>hbase.cluster.distributed</code>: Set to <code>true</code> for distributed mode.

HBase Java API

```

import org.apache.hadoop.conf.Configuration;
import org.apache.hadoop.hbase.HBaseConfiguration;
import org.apache.hadoop.hbase.client.Connection;
import org.apache.hadoop.hbase.client.ConnectionFactory;

Configuration config = HBaseConfiguration.create();
config.set("hbase.zookeeper.quorum", "localhost"); // Replace with your
ZooKeeper quorum

try (Connection connection = ConnectionFactory.createConnection(config)) {
    // Use the connection
} catch (Exception e) {
    e.printStackTrace();
}

```

```

import org.apache.hadoop.hbase.TableName;
import org.apache.hadoop.hbase.client.*;
import org.apache.hadoop.hbase.util.Bytes;

TableName tableName = TableName.valueOf("mytable");

try (Table table = connection.getTable(tableName)) {
    // Put data
    Put put = new Put(Bytes.toBytes("row1"));
    put.addColumn(Bytes.toBytes("cf1"), Bytes.toBytes("qual1"),
Bytes.toBytes("value1"));
    table.put(put);

    // Get data
    Get get = new Get(Bytes.toBytes("row1"));
    Result result = table.get(get);
    byte[] value = result.getValue(Bytes.toBytes("cf1"),
Bytes.toBytes("qual1"));
    String valueStr = Bytes.toString(value);
    System.out.println("Value: " + valueStr);

    // Scan data
    Scan scan = new Scan();
    try (ResultScanner scanner = table.getScanner(scan)) {
        for (Result row : scanner) {
            // Process each row
        }
    }
} catch (Exception e) {
    e.printStackTrace();
}

```