

# Perform Data Engineering on Microsoft HDInsight

The main purpose of the course is to give participants the ability to plan and implement big data workflows on HDInsight

## Who Should Attend

This primary audience for this course is data engineers, data architects, data scientists, and data developers who plan to implement big data engineering workflows on HDInsight

## Course Prerequisites

**Before attending this course, participants should have:**

- > Programming experience using R, and familiarity with common R packages
- > Knowledge of common statistical methods and data analysis best practices
- > Basic knowledge of the Microsoft Windows operating system and its core functionality
- > Working knowledge of relational databases

## Course Objectives

**After completing this course, participants will be able to:**

- > Deploy HDInsight Clusters
- > Authorising Users to Access Resources
- > Loading Data into HDInsight
- > Troubleshooting HDInsight
- > Implement Batch Solutions
- > Design Batch ETL Solutions for Big Data with Spark
- > Analyse Data with Spark SQL
- > Analyse Data with Hive and Phoenix
- > Describe Stream Analytics
- > Implement Spark Streaming Using the DStream API
- > Develop Big Data Real-Time Processing Solutions with Apache Storm
- > Build Solutions that use Kafka and HBase
- > Create Spark Streaming Applications

**Course Code:** 20775a  
**Course Duration:** 5 Days  
**SATV:** Yes

## Course Summary

**Module 1:** Getting Started with HDInsight  
**Module 2:** Deploying HDInsight Clusters  
**Module 3:** Authorising Users to Access Resources  
**Module 4:** Loading Data into HDInsight  
**Module 5:** Troubleshooting HDInsight  
**Module 6:** Implementing Batch Solutions  
**Module 7:** Design Batch ETL Solutions for Big Data with Spark  
**Module 8:** Analyse Data with Spark SQL  
**Module 9:** Analyse Data with Hive and Phoenix  
**Module 10:** Stream Analytics  
**Module 11:** Implementing Streaming Solutions with Kafka and HBase  
**Module 12:** Develop Big Data Real-Time Processing Solutions with Apache Storm  
**Module 13:** Create Spark Streaming Applications



Gold Data Analytics  
Gold Data Platform  
Silver Cloud Platform  
Silver Learning

register for this course today as places are strictly limited

 1300 927 394



 [contact@wardyit.com](mailto:contact@wardyit.com)

# Perform Data Engineering on Microsoft HDInsight

## Course Outline

### Module 1: Getting Started with HDInsight

This module introduces Hadoop, the MapReduce paradigm, and HDInsight

- > What is Big Data?
- > Introduction to Hadoop
- > Working with MapReduce Function
- > Introducing HDInsight

### Module 2: Deploying HDInsight Clusters

This module provides an overview of the Microsoft Azure HDInsight cluster types, in addition to the creation and maintenance of the HDInsight clusters. The module also demonstrates how to customize clusters by using script actions through the Azure Portal, Azure PowerShell, and the Azure command-line interface (CLI). This module includes labs that provide the steps to deploy and manage the clusters.

- > Identifying HDInsight Cluster Types
- > Managing HDInsight Clusters by using the Azure portal
- > Managing HDInsight clusters by using Azure PowerShell

### Module 3: Authorising Users to Access Resources

This module provides an overview of non-domain and domain-joined Microsoft HDInsight clusters, in addition to the creation and configuration of domain-joined HDInsight clusters. The module also demonstrates how to manage domain-joined clusters using the Ambari management UI and the Ranger Admin UI. This module includes the labs that will provide the steps to create and manage domain-joined clusters

- > Non-domain Joined Clusters
- > Configuring domain-joined HDInsight Clusters
- > Manage domain-joined HDInsight Clusters

### Module 4: Loading Data into HDInsight

This module provides an introduction to loading data into Microsoft Azure Blob storage and Microsoft Azure Data Lake storage. At the end of this lesson, you will know how to use multiple tools to transfer data to an HDInsight cluster. You will also learn how to load and transform data to decrease your query run time.

- > Storing Data for HDInsight Processing
- > Using Data Loading Tools
- > Maximising Value from Stored Data

### Module 5: Troubleshooting HDInsight

In this module you will learn how to interpret logs associated with the various services of Microsoft Azure HDInsight cluster to troubleshoot any issues you might have with these services. You will also learn about Operations Management Suite (OMS) and its capabilities.

- > Analyse HDInsight Logs
- > YARN Logs
- > Heat Dumps
- > Operations Management Suite

### Module 6: Implementing Batch Solutions

In this module, you will look at implementing batch solutions in Microsoft Azure HDInsight by using Hive and Pig. You will also discuss the approaches for data pipeline operationalization that are available for big data workloads on an HDInsight stack.

- > Apache Hive Solutions
- > HDInsight Data Queries using Hive and Pig
- > Operationalise HDInsight

### Module 7: Design Batch ETL Solutions for Big Data with Spark

This module provides an overview of Apache Spark, describing its main characteristics and key features. Before you start, it's helpful to understand the basic architecture of Apache Spark and the different components that are available. The module also explains how to design batch Extract, Transform, Load (ETL) solutions for big data with Spark on HDInsight. The final lesson includes some guidelines to improve Spark performance.

- > What is Spark?
- > ETL with Spark
- > Spark Performance

register for this course today as places are strictly limited

# Perform Data Engineering on Microsoft HDInsight

## Module 8: Analyse Data with Spark SQL

This module describes how to analyze data by using Spark SQL. In it, you will be able to explain the differences between RDD, Datasets and Dataframes, identify the uses cases between Iterative and Interactive queries, and describe best practices for Caching, Partitioning and Persistence. You will also look at how to use Apache Zeppelin and Jupyter notebooks, carry out exploratory data analysis, then submit Spark jobs remotely to a Spark cluster.

- > Implementing Iterative and Interactive Queries
- > Perform Exploratory Data Analysis

## Module 9: Analyse Data with Hive and Phoenix

In this module you will learn about running interactive queries using Interactive Hive (also known as Hive LLAP or Live Long and Process) and Apache Phoenix. You will also learn about the various aspects of running interactive queries using Apache Phoenix with HBase as the underlying query engine.

- > Implement Interactive Queries for Big Data with Interactive Hive
- > Perform Exploratory Data Analysis by using Hive
- > Perform Interactive Processing by using Apache Phoenix

## Module 10: Steam Analytics

The Microsoft Azure Stream Analytics service has some built-in features and capabilities that make it as easy to use as a flexible stream processing service in the cloud. You will see that there are a number of advantages to using Stream Analytics for your streaming solutions, which you will discuss in more detail. You will also compare features of Stream Analytics to other services available within the Microsoft Azure HDInsight stack, such as Apache Storm. You will learn how to deploy a Stream Analytics job, connect it to the Microsoft Azure Event Hub to ingest real-time data, and execute a Stream Analytics query to gain low-latency insights. After that, you will learn how Stream Analytics jobs can be monitored when deployed and used in production settings

- > Stream Analytics
- > Process Streaming Data from Stream Analytics
- > Managing Stream Analytics Jobs

## Module 11: Implementing Streaming Solutions with Kafka and HBase

In this module, you will learn how to use Kafka to build streaming solutions. You will also see how to use Kafka to persist data to HDFS by using Apache HBase, and then query this data

- > Building and Deploying a Kafka Cluster
- > Publishing, Consuming, and Processing Data using the Kafka Cluster
- > Using HBase to Store and Query Data

## Module 12: Develop Big Data Real-Time Processing Solutions with Apache Storm

This module explains how to develop big data real-time processing solutions with Apache Storm

- > Persist Long Term Data
- > Stream Data with Storm
- > Create Storm Topologies
- > Configure Apache Storm

## Module 13: Using R Services with Machine Learning

This module describes how to use R and R server with Azure machine learning, and explain how to deploy and configure SQL Server and support R services

- > R and R Server overview
- > Using R server with machine learning
- > Using R with SQL Server

register for this course today as places are strictly limited

 1300 927 394



 [contact@wardyit.com](mailto:contact@wardyit.com)