



Hortonworks Data Platform Developer: Custom YARN Applications H7G72S

This 2-day hands-on training course teaches students how to develop custom YARN applications for Apache Hadoop. Students will learn the details of the YARN architecture, the steps involved in writing a YARN application, the details of writing a YARN client and ApplicationMaster, and how to launch Containers. Applications are developed using Eclipse and Gradle connected remotely to a 7-node HDP 2.1 cluster running in a virtual machine that the students can keep for use after the training.

Hortonworks Data Platform Developer: Custom YARN Applications

Price USD \$1,800

**Links to local
schedules,
pricing and
registration** [US/Canada](#)
[Mexico/Latin America](#)
[Brazil](#)

HP course # H7G72S

Category Big Data

Duration 2 days

Audience

- Software engineers familiar with Java who need to develop YARN applications on Hadoop 2.x by writing custom YARN clients and ApplicationMasters in Java

Prerequisites

Experienced Java developers who have:

- Attended the H7G67S Hortonworks Data Platform Developer: Java course; or
- attended the Hortonworks Data Platform Developer: Apache Pig and Hive course; or
- possess similar Hadoop development knowledge and understand HDFS and the MapReduce framework

Course objectives

At the completion of the course students will be able to:

- Explain the architecture of YARN
- Explain the lifecycle of a YARN application
- Write a YARN client application
- Run a YARN application on a Hadoop 2.x Cluster
- Monitor the status of a running YARN application

- View the aggregated logs of a YARN application
- Configure a ContainerLaunchContext
- Define a LocalResource for sharing application files across the cluster
- Write a YARN ApplicationMaster
- Explain the differences between synchronous and asynchronous ApplicationMasters
- Allocate Containers in a cluster
- Launch Containers on NodeManagers
- Write a custom Container to perform specific business logic
- Explain the job schedulers of the ResourceManager
- Define queues for the Capacity Scheduler

Benefits to you

- This course will provide in depth explanation on how to develop custom YARN applications

Course outline

Day 1

- Unit 1: The YARN Architecture
- Unit 2: Overview of a YARN Application
- Unit 3: Writing a YARN Client

Day 2

- Unit 4: Writing a YARN ApplicationMaster
- Unit 5: Containers
- Unit 6: Job Scheduling

Hands-On Labs

Students will work through the following lab exercises using the Hortonworks Data Platform 2.1:

- Running a YARN Application
- Setup a YARN Development Environment
- Writing a YARN Client
- Submitting an ApplicationMaster
- Writing an ApplicationMaster
- Requesting Containers
- Running Containers
- Writing Custom Containers

Learn more at

hpe.com/us/training/bigdata