



Hortonworks Data Platform Developer—Custom YARN Applications (EDU-PRV-DEV- YARN-200) H7G72S

HPE course number	H7G72S
Course length	2 days
Delivery mode	ILT
View schedule, local pricing, and register	View now
View related courses	View now

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.

Why HPE Education Services?

- IDC MarketScape leader 4 years running for IT education and training*
- Recognized by IDC for leading with global coverage, unmatched technical expertise, and targeted education consulting services*
- Key partnerships with industry leaders OpenStack®, VMware®, Linux®, Microsoft®, ITIL, PMI, CSA, and (ISC)²
- Complete continuum of training delivery options—self-paced eLearning, custom education consulting, traditional classroom, video on-demand instruction, live virtual instructor-led with hands-on lab, dedicated onsite training
- Simplified purchase option with HPE Training Credits

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

*Realize Technology Value with Training, IDC Infographic 2037, Sponsored by HPE, January 2016

- 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

Detailed 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/ww/en/training/contactus.html

Follow us:



© Copyright 2015–2016 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft is either a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries. Java is a registered trademark of Oracle and/or its affiliates. The OpenStack Word Mark is either a registered trademark/service mark or trademark/service mark of the OpenStack Foundation, in the United States and other countries and is used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation or the OpenStack community. Pivotal and Cloud Foundry are trademarks and/or registered trademarks of Pivotal Software, Inc. in the United States and/or other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other third-party trademark(s) is/are property of their respective owner(s).

c04577750, December 2016, Rev. 1