Hadoop On Demand

Table of contents

1. Introduction

Hadoop On Demand (HOD) is a system for provisioning virtual Hadoop clusters over a large physical cluster. It uses the Torque resource manager to do node allocation. On the allocated nodes, it can start Hadoop Map/Reduce and HDFS daemons. It automatically generates the appropriate configuration files (hadoop-site.xml) for the Hadoop daemons and client. HOD also has the capability to distribute Hadoop to the nodes in the virtual cluster that it allocates. In short, HOD makes it easy for administrators and users to quickly setup and use Hadoop. It is also a very useful tool for Hadoop developers and testers who need to share a physical cluster for testing their own Hadoop versions.

2. Documentation

Please go through the following to know more about using HOD

- [Hod Admin Guide](#) : This guide will walk you through an overview of architecture of HOD, prerequisites, installing various components and dependent software, and configuring HOD to get it up and running.
- [Hod User Guide](#) : This guide will let you know about how to get started on running hod, its various features, command line options and help on troubleshooting in detail.
- [Hod Configuration Guide](#) : This guide discusses about onfiguring HOD, describing various configuration sections, parameters and their purpose in detail.