

## Apache Ignite

### Course Summary

#### Description

Apache Ignite is an open-source distributed database, caching and processing platform designed to store and compute on large volumes of data across a cluster of nodes.

This course offers an introduction to Apache Ignite. It is intended for data scientists and software engineers. It gives them practical level of experience, achieved through a combination of about 50% lecture, 50% lab work.

#### Topics

- Overview
- In-memory data grid
- Highly available service grid
- Streaming (event processing for IoT and fast data)
- In-memory computing use cases for high-performance
- Use cases
- High volume of ACID transactions in your system
- Database bottleneck in your application
- Developing and deploying microservices in a distributed fashion
- Improving the Hadoop ecosystem (OLAP) without making changes in your existing MapReduce jobs
- Sharing Spark RDD directly in-memory (without storing the state into the disk)
- Processing continuous never-ending streams and complex events of data.
- Running distributed computations in parallel fashion to gain high performance

#### Audience

This course is designed for data scientists and developers.

#### Prerequisites

Before taking this course, students should have the following skills and experience:

- Familiarity with Big Data
- Be able to navigate Linux command line
- Basic knowledge of command line Linux editors (VI / nano)

#### Duration

Three days