

Serverless Technologies (Lambda and DynamoDB)

Course Summary

Description

Serverless Technologies Is a design principle and best practice in cloud computing. It teaches one to design solutions that do not depend on your running specific servers. Rather, the cloud decides which server and how will implement your functionality.

Two of the enabling technologies for serverless computing are AWS Lambda and AWS DynamoDB. Lambda allows the developer to specify the computing functions to be performed, leaving the implementation and scalability, as well as the SLA, to the cloud. DynamoDB is Amazon's NoSQL database that runs as a service, making it a perfect fit for Lambda.

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

Audience

This course is designed for Data Scientists and Developers

Topics

- Serverless Computing Introduction
- AWS Lambda
- AWS DynamoDB

Prerequisites

- General familiarity with cloud computing; if required, an overview will be provided
- Be able to navigate the linux command line, although most labs are in the browser
- Students need to come with their own AWS account. If needed, an arrangement for the students' account can be provided.

Duration

Three Days

Serverless Technologies (Lambda and DynamoDB)

Course Outline

- I. **Serverless Computing Introduction**
 - A. What is serverless computing?
 - B. Why use serverless computing?
 - C. What are the benefits of serverless computing?
 - D. AWS Serverless Overview
 1. Cloud logic layer
 2. Orchestration and state management
 3. Data sources
 4. Application modeling framework
 5. Developer ecosystem
 6. Application and integration library
 7. Security and access control
 8. Reliability and performance
 9. Global scale and reach
- II. **AWS Lambda**
 - A. Benefits of AWS Lambda
 - B. Lambda architecture
 - C. Lambda use cases
 - D. Writing code for Lambda
 - E. Managing code for Lambda
 - F. What can be done from Lambda
 - G. Data sources
 - H. Scalability advice and best practices
- III. **AWS DynamoDB**
 - A. Intro NoSQL
 1. NoSQL overview
 2. CAP theorem
 3. When is NoSQL appropriate
 4. Columnar storage
 5. NoSQL Ecosystem
 - B. AWS NoSQL (DynamoDB) in depth
 1. Amazon DynamoDB Accelerator
 2. Key-value and document data model support
 3. Develop locally on your desktop
 4. Secondary indexes
 5. Streams
 6. Triggers
 - C. NoSQL and DynamoDB in the real world
 1. Creating and using secondary indexes
 2. Composite keys
 3. Time series data
 4. Best practices for time series data
 5. Counters
 6. Data Modeling Labs: Group design session
 - Multiple use cases from various domains are presented
 - Students work in groups to come up with designs and models
 - Discuss various designs, analyze decisions
 - Lab: implement one of the scenarios