Course Outline

ProTech Professional Technical Services, Inc.

AWS Security, Big Data, and Lambda Fundamentals

Course Summary

Description

In this course, you will learn the fundamentals of AWS Security, EMR Big Data and AWS Serverless Framework. You will be introduced to the concept of shared-responsibility model to understand the role of AWS and your role in securing your Serverless applications. You will learn the best practices in designing highly redundant applications across Availability Zones utilizing Elastic Load Balancers and Autoscaling. You will perform hands-on exercises throughout the course to reinforce the lessons being taught

Objectives

Attendees will leave with an understanding of AWS Serverless, specifically Lambda and how to use it with serverless software frameworks. You will be able to build and deploy serverless applications with testing instrumentation, leverage multiple serverless technologies including static hosted websites and backend data stores for a complete end to end serverless product creation and management.

Attendees will also leave with an understanding of AWS Security Services including a newly announced product Security Hub. Attendees will leave with an understanding of AWS Big Data analytics with Amazon Elastic MapReduce and will learn how to write lambda functions connecting to EMR clusters. Recent announcements of services in Big Data analytics such as Lake Formation will be discussed.

Topics

- Introduction to AWS Cloud
- AWS Services Overview
- IAM (Identity and Access Management)
- CloudTrail
- AWS Security Hub
- Virtual Private Cloud (VPC) and Security
- AWS Big Data with Elastic MapReduce
- AWS Lake Formation
- What is Function-as-a-Service/Serverless?
- API/Microservices routing and triggering
- Example Application 1 Overview: Static Blog with Hugo and Lambda
- Adding a comments section to our static site
- Architectural Review
- Gotchas and limitations of FaaS
- Serverless Frameworks
- Debugging and tracing in Lambda
- Testing lambda functions
- Canary Deployments
- Deploying an application with an automated canary
- Architectural Review

Audience and Prerequisites

The audience is any intermediate or advanced developer with a general concept of cloud computing, virtualization, and ideally at least an understanding of what containers are.

Knowledge of automated testing, CI/CD, and pipeline instrumentation would be beneficial.

Duration

Two days

Course Outline

AWS Security, Big Data, and Lambda Fundamentals

Course Outline

I. *Introduction to AWS Cloud*
   A. The AWS Global Infrastructure: Regions and Availability Zones
   B. Designing for AWS cloud: The well-architected guidelines

II. *AWS Services Overview*
   A. The Management console and CLI
   B. Different Categories of AWS services
   C. Review of the Security and Serverless services
   D. Review of Elastic MapReduce services

III. *IAM (Identity and Access Management)*
   A. The root and IAM users
   B. Creating groups, users
   C. Protecting the IAM Account: Multi-Factor Auth, Password Policies
   D. Passwords, Keys, Roles

IV. *CloudTrail*
   A. Introduction to CloudTrail
   B. Configuring CloudTrail
   C. Connecting CloudTrail to AWS CloudWatch
   D. CloudTrail and Lambda functions

V. *AWS Security Hub*
   Introduction to Security Hub

VI. *Virtual Private Cloud (VPC) and Security*
   A. Creating VPCs, subnets, IGW, and Route Tables
   B. Using Security Groups to protect your resources
   C. Network Access Control List (NACL) as Extra layer of defense

VII. *AWS Big Data with Elastic MapReduce*
   A. Introduction to AWS Big Data Services
   B. Creating an Elastic MapReduce cluster
   C. Performing MapReduce on AWS
   D. Using Elastic MapReduce utilities

VIII. *AWS Lake Formation*
   Introduction to AWS Lake Formation

IX. *What is Function-as-a-Service/Serverless?*
   A. Serverless is containers
   B. What are containers? (briefly)

X. *API/Microservices routing and triggering*
   A. API Gateway overview
   B. Message Queues
   C. DynamoDB Triggers
   D. SQS Triggers
   E. S3 Bucket Triggers

XI. *Example application 1 overview: Static blog with Hugo and Lambda*
   A. What is Hugo (Static site generator)
   B. How to host static files
   C. What is GitHub / Creating a GitHub account or using an existing one

XII. *Adding a comments section to our static site*
   A. LAB: Use a Lambda function to save and fetch data from a DynamoDB table
   B. LAB: Integrate the dynamic Lambda function content with a static site hosted on s3

XIII. *Architectural Review*
   A. Review used concepts of front/backend development with static front ends and FaaS backends
   B. Review used concepts of DevOps/Pipeline management in FaaS
   C. Stream processing / Functional programming vs. OOP – How to properly leverage FaaS
   D. Scheduled Tasks

XIV. *Gotchas and Limitations of FaaS*
   A. Billing details
   B. Error tracing and failure scenarios
   C. Thread pools and concurrency limitations
   D. Cold/Warm/Hot starts and how to mitigate or manage them
   E. Timeouts and service thresholds

Course Outline (con't)

XV. Serverless Frameworks

A. Serverless framework overview
B. AWS SAM overview
C. LAB: Deploy a simple Lambda based demo application using Serverless
D. LAB: Deploy a simple Lambda based demo application user AWS SAM

XVI. Debugging and tracing in Lambda

A. Local debugging using AWS SAM
B. Remote tracing live Lambda functions

XVII. Testing lambda functions

A. LAB: Local Lambda unit testing using AWS SAM
B. LAB: Local Lambda unit testing with Docker-Lambda
C. LAB: Local Lambda integration testing using AWS SAM

XVIII. Canary Deployments

A. What canary deployments are and how they are used
B. How to version Lambda functions and enable weighted routing
C. LAB: Manually tag a Lambda function with a revision alias and apply weighted routing
D. LAB: Create an API gateway with API versions calling the appropriate Lambda function

XIX. Deploying an application with an automated canary

A. LAB: Deploy an application from a CI/CD pipeline that automatically performs a canary rollout with a versioned API
B. Use a Lambda function on a scheduled task to monitor Cloudwatch events
C. Interact with the API gateway via Lambda to dynamically increase the routing weight
D. Backout the deployment if an error threshold is hit and send a notification

XX. Architectural Review

A. How large should functions be?
B. Advantages and disadvantages of having multiple functions in a single function with multiple entry points
C. Live monitoring and alerting for Lambda functions