aws academy

Cloud Computing Architecture: Curriculum Objective Details

Below are estimated timings to deliver each component of the Cloud Computing Architecture curriculum.
Please note that actual delivery times will vary from class to class.

Unit 1 : Introduction to Cloud Computing

Unit 1.01	Lecture	What is Cloud Computing?	35 mins
Unit 1.01	**Knowledge Assessment**	What is Cloud Computing?	10 mins

Lecture and Assessment: This module introduces students to concepts of cloud computing, a comparison of cloud and on-premises computing, and a description of the AWS Global Infrastructure.
At the end of this module, the student will be able to:

- Define cloud computing
- Identify what purposes cloud computing can be used
- Compare cloud computing to on-premises computing
- Distinguish between AWS regions and availability zones
- Describe the purpose of AWS edge locations
- Categorize AWS services as infrastructure, foundation services, platform services, or applications

Unit 1.02	Lecture	Leveraging Cloud Computing	75 mins
Unit 1.02	**Knowledge Assessment**	Leveraging Cloud Computing	10 mins

Lecture and Assessment: This module shows the learner what components are used to create a simple implementation for a single user, and increases complexity adding cloud services to address high availability, and multi-tier architectures. Auto Scaling is introduced showing the advantage of automated provisioning in response to actual demand.
At the end of this module, the student will be able to:

- Describe a simple AWS architecture:
  - Generally characterize the use of EC2 in a cloud computing solution
  - Identify how an AMI is used
  - Generally characterize the use of Route 53 in a cloud computing solution
- Describe a simple high availability architecture
  - Describe how a multi-AZ approach provide high availability
  - Describe how Elastic Load Balancer works
- Describe the use of S3 and CloudFront to improve efficiency
- Describe how to use DynamoDB or ElastiCache to achieve a stateless tier
  - Explain the benefit of this approach when scaling compute instances
- Describe how Auto Scaling works
  - Generally characterize the use of CloudWatch in Auto Scaling

Unit 1.03	Lecture	Cloud Economics and Total Cost of Ownership	60 mins
Unit 1.03	**Knowledge Assessment**	Cloud Economics and Total Cost of Ownership	10 mins

Lecture and Assessment: This module explains cost optimization and total cost of ownership (TCO). Cost optimization discusses right-sizing instance types for greater efficiency; leveraging reserved instances; increasing elasticity; and tools for monitoring and analyzing a system over time for continual optimization. TCO describes what elements to consider in a comparison analysis of acquisition and maintenance costs for an on-premises system and a cloud-based system.
At the end of this module, the student will be able to:

- Identify the four pillars of cost optimization:
  - Analyze CPU, RAM, storage, and network utilization to identify the most cost-effective instance type
  - Evaluate under what circumstances reserved instances (RIs) provide the best coverage at the most cost-effective point
  - Describe how elasticity and auto scaling can drive cost savings
  - Describe opportunities that could be discovered by collecting and analyzing system data
- Define what is meant by Total Cost of Ownership
  - Identify categorically the breadth of costs associated with
- operating and maintaining an on-premises computing system
  - Identify advantages achieved by operating the same system in a cloud platform
  - Identify AWS tools available to assist in cost evaluation

Unit 2 : Getting Started with AWS

Unit 2.01	Lecture	AWS Compute, Storage & Networking	75 mins
1	Lab	Creating Amazon EC2 instances with Microsoft Windows	90 mins
2	Lab	Build your Virtual Private Cloud and Launch a Web Server	90 mins
3	Lab	Working with Amazon Elastic Block Store	30 mins
Unit 2.01	**Knowledge Assessment**	AWS Compute, Storage & Networking	10 mins

Lecture and Assessment: This module goes into some of the foundational services of AWS in Compute (EC2), Networking (VPC), and Storage (S3, EBS). Three labs are included in this module.
At the end of this module, the student will be able to:

- Understand EC2 concepts (instances vs. servers, instance types and families, ephemeral vs. persistent storage, AMIs, instance metadata and user data)
  - Describe how to use an AMI to launch an EC2
  - Identify optimal EC2 instance types and purchase options based on unique requirements
  - Describe metadata best practices and access
- Understand VPC concepts (subnets, security, networking, VPN).
  - Describe how security groups are used
- Understand AWS storage concepts, features and limits including:
  - S3 (requests, buckets, objects, access, protecting data, replication, request routing, optimization, lifecycle management with Glacier
  - EBS (volumes,snapshots,optimization,encryption,performance)
  - EC2 Instance Store (SSD, swap volumes, optimizing disk performance)
  - EC2 Instance Store (SSD, swap volumes, optimizing disk performance)

Unit 2 : Getting Started with AWS (continued)

Lab 1: By the end of this lab, you will be able to:
- Create a new Amazon EC2 instance with Windows Server 2012 R2 & IIS.
- Create a security group to restrict access to the server's resources.
- Launch the instance.
- Access the instance.
- Associate an Elastic IP address with your Amazon EC2 instance.

Lab 2: After completing this lab, you will be able to:
- Create a VPC.
- Create subnets.
- Configure a security group.
- Launch an EC2 instance into a VPC.

Lab 3: By the end of this lab, you will be able to:
- Create an Amazon EBS volume.
- Attach the volume to an instance.
- Configure the instance to use the virtual disk.
- Create an Amazon EBS snapshot.
- Restore the snapshot.

Unit 2.02	Lecture	AWS Security, Identify and Access Management (IAM)	40 mins
4	Lab	Introduction to AWS Identify and Access Management (IAM)	60 mins
Unit 2.02	**Knowledge Assessment**	AWS Security, Identify and Access Management (IAM)	10 mins

Lecture and Assessment: This module introduces security concepts, AWS security concepts, and IAM. The module is followed by an instructor-led demo of IAM.
At the end of this module, the student will be able to:
- Explain the AWS Shared Responsibility Model.
- Describe the security measures AWS provides.
  - Describe SSL endpoints
  - Describe how security groups are used
- Understand IAM concepts (users, groups, roles, policies).
  - Differentiate between authorized and authenticated
  - Describe what the AWS policy simulator is used for
  - Identify use cases, features and limits for AWS Security Token Service (STS)
  - Differentiate between AWS services that do or do not support

Lab 4: By the end of this lab, you will be able to:
- Exploring pre-created IAM users and groups
- Inspecting IAM policies as applied to the pre-created groups
- Following a real-world scenario, adding users to groups with specific capabilities enabled
- Updating passwords for users.
- Locating and using the IAM sign-in URL
- Experimenting with the effects of policies on service access

Unit 2.03	Lecture	AWS Database Options	45 mins
5	Lab	Build your Database Server and interact with your Database using an Application	90 mins
Unit 2.03	**Knowledge Assessment**	AWS Database Options	10 mins

Lecture and Assessment: This module introduces database concepts and dives into AWS databases including RDS and DynamoDB.
At the end of this module, the student will be able to:
- Differentiate between NoSQL and. SQL database models.
  - Determine factors that affect what database type would be most effective
  - Differentiate between AWS managed and self-managed services
- Understand RDS concepts (DB instances, security groups, DB parameter groups, DB option groups, RDS interfaces, supported Db applications)
  - Describe methods of accessing RDS
  - Describe the RDS backup process and lifecycle
- Understand DynamoDB concepts (data model, supported operations, provisioned throughput, accessing DynamoDB).

Lab 5: After completing this lab, you will be able to:
- Launch an Amazon RDS DB instance with high availability.
- Configure the DB instance to permit connections from your web server.
- Open a web application and interact with your database.

Unit 2.04	Lecture	Elasticity and Management Tools	30 mins
6	Lab	Scale and Load Balance your Architecture	110 mins
Unit 2.04	**Knowledge Assessment**	Elasticity and Management Tools	10 mins

Lecture and Assessment: This module introduces the "trio of services" -- Auto Scaling, Elastic Load Balancing, and CloudWatch. The module also introduces Trusted Advisor as a management tool.
At the end of this module, the student will be able to:
- Understand Auto Scaling concepts (launch configurations, auto scaling groups, scaling plans, auto scaling lifecycle).
- Understand ELB concepts (request routing, load balancer types, back-end instances, listeners).
  - Differentiate between classic and application load balancers
- Understand AWS Management tools:
  - Describe CloudWatch concepts, architecture, and what metrics it monitors and reports on:
    - o Differentiate between basic and detailed monitoring
    - o Describe use cases that leverage CloudWatch with Auto Scaling
    - o Identify different methods for accessing CloudWatch
    - o Explain how CloudWatch alarms work and what AWS services can be monitored
  - Describe Trusted Advisor concepts and use cases for optimizing cost, security, fault tolerance, and performance improvement
    - o Interpret status check color codes in the Trusted Advisor dashboard

Lab 6: After completing this lab, you will be able to:
- Create an Amazon Machine Image (AMI) from a running instance.
- Create a load balancer.
- Create a launch configuration and an Auto Scaling group.
- Autoscale new instances within a private subnet.

Unit 3 : Architecting on AWS

Part 1: Introduction to System Design

Unit 3.01	Lecture	AWS Essentials Review	30 mins

Lecture: This module reviews the content that was provided in Units 1 and 2. Students should have already taken the knowledge assessments for all of the modules in Units 1 and 2. This review is intended as a classroom discussion both as a refresher, and to help the instructor identify any learning gaps before continuing. The Instructor is responsible for gauging the level of knowledge of the students and determining if a deeper discussion is necessary to ensure that all students grasp the basics of these services. There is no assessment for this module.

Unit 3.02	Lecture	Designing Your Environment	45 mins
7	Lab	Deploy a Web Application onto AWS	180 mins
Unit 3.02	**Knowledge Assessment**	Designing Your Environment	10 mins

Lecture and Assessment: This module provides insight on system design based on best practices:

- Design an AWS environment based on guidelines and patterns for Region and AZ selection.
  - Describe the concept data sovereignty and how that may impact choices for system design.
  - Describe system design for high availability and fault tolerance.
- Describe system design for Multi- Account patterns, Multi-VPC patterns, and Subnet structures
  - Describe use cases for multi-account and multi-VPC architectures
  - Describe features and limitations of multi-account and multi-VPC architectures
  - Identify how rout table are used with a VPC and subnet
- Describe IP subnetting including CIDR notation, and IP addresses
  - Differentiate between public and private subnets
  - Calculate the number of IP addresses associated with CIDR notation
  - Determine CIDR block sizes for a VPC and related subnets
  - Analyze the number of hosts in a subnet, and number of IPs provided in a public or private subnet
  - Evaluate use cases for fewer subnets of a larger size with more subnets of a smaller size
  - Describe methods for securing hosts within a VPC
- Describe how security groups work with transfer protocols, subnets and AZs
  - Describe features and concepts of Internet gateways and their impact on access within subnets
  - Distinguish VPC NAT gateways
  - Describe capabilities and limits of VPC peering

Lab 7: After completing Lab 7, you will be able to:
- Create an IAM User and attach a permission policy for controlling access to services
- Create a Virtual Private Cloud (VPC) with an Internet Gateway (IGW) and a Public Subnet
- Create an S3 Bucket and upload static objects into the bucket
- Create a DynamoDB table and populate with items
- Deploy a web application in a cost optimized and scalable manner utilizing EC2, S3, and DynamoDB

1	**Discussion**	Forklift an Existing Application onto AWS	30 mins

The goal of this group discussion is to summarize what was covered in Part 1 of Unit 3 and apply the concepts to the architecture.

Unit 3.03	Lecture	System Design for High Availability (Part I & II)	45 mins
8	Lab	Making your Environment Highly Available	150 mins
Unit 3.03	**Knowledge Assessment**	System Design for High Availability (Part I & II)	10 mins

Lecture and Assessment: This module emphasizes the best practice of "Avoid Single Points of Failure" and explains the concepts of High Availability and Fault Tolerance.

- Define 'high availability' and 'single point of failure'.
  - Differentiate between RTO and RPO
  - Compare fault tolerance with high availability
  - Design highly available and fault tolerant systems leveraging AZs.
  - Improve system availability using AWS services that are inherently highly available.
- Describe the functionality of ELB (health checks, connection draining, cross-zone load balancing)
  - Explain how ELB responds if an EC2 instance were to fail
  - Identify how an ELB relates to a VPC and subnet
- Describe the functionality and features of Route 53
  - Compare simple routing with geolocation routing and weighted round robin
  - Define what is meant by an elastic IP address
- Describe the capability and limitations of AWS Direct Connect

Lab 8: After completing this lab, you will be able to:
- Create an image of an existing Amazon EC2 instance and use it to launch a new instance.
- Create an Amazon ELB load balancer and attach it to Amazon EC2 instances.
- Create an AWS NAT Gateway.
- Create private subnets and launch Amazon EC2 instances into them.
- Edit private subnet route tables and security groups to intelligently control access.
- Test an AWS NAT Gateway.

Unit 3 : Architecting on AWS (continued)

Part 2: Automation and Serverless Architectures

Unit 3.04	Lecture	Event Driven Scaling	60 mins
9	Lab	Using Auto Scaling with AWS Lambda and Lifecycle Hooks	125 mins
Unit 3.04	**Knowledge Assessment**	Event Driven Scaling	10 mins

Lecture and Assessment: This module discusses enabling scalability based on workload needs and requirements.

- Differentiate between vertical and horizontal scaling
- Identify the functionality and capabilities of CloudWatch
- Explain how CloudWatch works with Auto Scaling using thresholds to scale horizontally.
  - Describe how the launch configuration is used to define new EC2 instances, AMIs, instance types, and security groups.
  - Identify what parameters are set in an Auto Scaling group
  - Differentiate between Auto Scaling groups, launch configurations, policies, scheduled actions, and lifecycle hooks.
  - Construct Auto Scaling policies based on unique application needs
- Compare EC2 Auto Recovery with Auto Scaling
- Explain "Scaling Datastores" concepts
  - Describe methods for scaling RDS either vertically or horizontally
  - Describe how putting a cache in front of your RDS instance may increase read speed
  - Describe what database sharding is and how it may improve availability and write performance
- Describe how AWS Lambda is used as a serverless architecture component
  - Identify compatible programming languages that can be used with Lambda
  - Describe a use case leveraging Lambda in an event-driven model

Lab 9: After completing this lab, you will be able to:
- Manually scale an Auto Scaling group.
- Implement an Auto Scaling lifecycle hook that invokes a Lambda function.

Unit 3.05	Lecture	Automating Your Infrastructure	60 mins
10	Lab	Creating an Amazon VPC with AWS CloudFormation	30 mins
Unit 3.05	**Knowledge Assessment**	Automating Your Infrastructure	10 mins

Lecture and Assessment: This module introduces the concept of "Infrastructure as Code" and discusses the need for automating the deployment of resources and environments to ensure consistency, repeatability, maintainability, and parallelization.

- Explain how the concept of 'infrastructure as code' relates to creating and managing your infrastructure
- Describe the functionality and implications of CloudFormation
  - Describe the use and format of templates (resources, resource groupings, metadata, description, parameters, mappings, conditions)
  - Identify the programming language used to create and manage a template
  - Describe what the CloudFormation Engine is
  - Identify what resources require a VPC gateway attachment before being created
  - Describe the functionality of CloudFormation assets (DependsOn attribute, UserData, wait condition, creation policy)
- Identify other resource deployment and configuration automation options offered by AWS

Lab 10: After completing this lab, you will be able to:
- Examining the sections and components of an AWS Cloud Formation template.
- Launching a stack using an AWS CloudFormation template.
- Modifying a stack's protection and deleting a stack.

1	**Project**	Mid-Curriculum Project	380 mins

After completing this project, you will be able to:

- Configuring access permissions to conform with AWS best practices
- Building networks that conform to AWS best practices while providing all the necessary network services to the applications in their different environments.
- Configuring auditing to track all user actions
- Building an architecture that matches the current architecture at the server hosting company and that can handle doubling the number of servers.

Unit 3.06	Lecture	Decoupling Your Infrastructure	60 mins
Unit 3.06	**Knowledge Assessment**	Decoupling Your Infrastructure	10 mins

Lecture and Assessment: This module introduces design patterns for decoupling and reducing inter-dependencies between tiers. It discusses best practices for using microservices and designing solutions with components.

- Describe best practices of loose coupling and the related use of microservices
  - Identify the advantages of loosely coupling your components; and designing services, not servers
  - Compare how to implement using SQS and SNS including parameters and limits of both
- Explain how to store and retrieve processing output using DynamoDB
- Explain how to use the API Gateway as a way of scaling API calls to access data, business logic, or functionality on back-end services
- Explain how to use AWS Lambda to execute code in response to events within milliseconds and without the customer having to manage any compute resources
- Demonstrate serverless approaches to decoupling with example architectures and use cases

Unit 3 : Architecting on AWS (continued)

Part 2: Automation and Serverless Architectures (continued)

Unit 3.07	Lecture	Designing Web-Scale Storage	45 mins
11	Lab	Caching Static Files with Amazon CloudFront	115 mins
12	Lab	Implementing a Serverless Architecture with AWS Managed Services	120 mins
Unit 3.07	**Knowledge Assessment**	Designing Web-Scale Storage	20 mins

Lecture and Assessment: This module demonstrates how to improve performance by optimize static content storage using S3, and to reduce latency by caching static content at Edge Locations using CloudFront. Further discussion provides guidance on data-store solutions and a comparison of Relation Database and NoSQL.

- Optimize a system design that leverages S3 for the storage of static content
- Describe how S3 handles data
  - Identify object naming conventions for S3
  - Describe the advantage of pre-pending a random hash key at the beginning of an object name to split the operations across separate partitions
- Explain how to use CloudFront to cache static content at Edge Locations
  - Describe advantages of caching (lower latency, increase performance, reduce costs)
  - Explain distribution management using CloudFront
- Evaluate the right choice for a data-store solution
- Determine when a Relation Database is needed and when NoSQL should be used.
  - Identify capabilities and limits of DynamoDB
  - Differentiate between features and use cases for DynamoDB and RDS
- Compare and differentiate Amazon Aurora with other RDS platform engines
  - Identify the capabilities and limits of Aurora

Lab 11: After completing this lab, you will be able to:
- Create and configure an Amazon CloudFront web distribution.
- Update and invalidate your content on Amazon CloudFront.
- Test your content from both Amazon S3 and Amazon CloudFront.

Lab 12: After completing this lab, you will be able to:
- Use AWS managed services to implement a serverless architecture.
- Set up Lambda functions to act as triggers in a DynamoDB table.

2	**Discussion**	Build an Environment for a New Web Application	45 mins

The goal of this group discussion is to apply the concepts learned during Part 2 of Unit 3, and design a new solution bridging off the Group Discussion 1 Architecture.

Part 3: Well-Architected Best Practices

Unit 3.08	Lecture	Introducing Well-Architected Framework	30 mins
Unit 3.08	**Knowledge Assessment**	Introducing Well-Architected Framework	10 mins

Lecture and Assessment: This module introduces the Well-Architected Framework and provides a high-level overview of each of the four pillars and the design principles that guide the Well-Architected Framework.

- Differentiate between a traditional environment and cloud environment
- Describe the advantage of automation in a cloud environment
- Identify the purpose of the Well-Architected Framework in evaluating architectures
- Identify the four security pillars

Unit 3.09	Lecture	Well-Architected Pillar 1: Security	75 mins
Unit 3.09	**Knowledge Assessment**	Well-Architected Pillar 1: Security	10 mins

Lecture and Assessment: This module provides greater depth in the Security pillar of the Well-Architected Framework.

- Describe how to secure data at every layer in the application (OS firewalls, security groups, network ACLs, route tables)
  - Describe measures for preventing DDoS attacks including services with built-in mitigation features
- Describe how to provide security by leveraging AWS tools and services: Amazon Inspector, CloudFront, AWS KMS, RDS Security, Identity Federation, DDoS Mitigation, EBS Encryption, S3 SSE, STS
  - Describe what Amazon Inspector is used for, and how it evaluates your applications
  - Describe how CloudFront can be used to protect data
  - Describe how to leverage subnets and security groups with EC2 to improve security
  - Describe the use of KMS master keys, encryption, and how AWS KMS compares with AWS CloudHSM
  - Describe how data is protected on Amazon S3 and Glacier
  - Identify types of security groups used with Amazon RDS, data at rest and data in transit
  - Identify options for running AWS Directory Service
  - Describe use cases for AWS Security Token Service (STS)

Unit 3 : Architecting on AWS (continued)

Part 3: Well-Architected Best Practices (continued)

Unit 3.10	Lecture	Well-Architected Pillar 2: Reliability	60 mins
13	Lab	Multi-Region Failover with Amazon Route 53	125 mins
Unit 3.10	**Knowledge Assessment**	Well-Architected Pillar 2: Reliability	10 mins

Lecture and Assessment: This module provides depth in the Reliability pillar of the Well-Architected Framework.

- Describe how to leverage AWS architectures to improve the reliability of a system
- Describe the application of AWS tools and services for increased reliability: IAM, VPC, AWS CloudTrail, AWS Config, AWS CloudFormation
- Describe architectural patterns for disaster recovery
  - Differentiate between RTO and RPO solutions
  - Compare recovery patterns: backup and restore, pilot light, low-capacity standby, and active-active
- Describe use cases for Amazon WorkSpace
- Identify which AWS services offer automated scaling, control access, supports change management and failure management

Lab 13: After completing this lab, you will be able to:
- Use Route 53 to configure cross-region failover of a web application.
- Use Route 53 health checks to determine the health of a resource.

Unit 3.11	Lecture	Well-Architected Pillar 3: Performance Efficiency	45 mins
Unit 3.11	**Knowledge Assessment**	Well-Architected Pillar 3: Performance Efficiency	10 mins

Lecture and Assessment: To provide depth in the Performance Efficiency pillar of the Well-Architected Framework.

- Describe how to tune or offload components of your system to improve the performance of a system.
- Describe the application of AWS tools and services for increased performance: EBS, S3, Glacier, Amazon RDS, Amazon DynamoDB, Amazon CloudFront, Auto-Scaling, ElastiCache, Kinesis Streams
  - Identify performance priorities for choosing an EC2 instance size
  - Evaluate storage and database solutions based on a range of hot to cold performance and durability

Unit 3.12	Lecture	Well-Architected Pillar 4: Cost Optimization	45 mins
Unit 3.12	**Knowledge Assessment**	Well-Architected Pillar 4: Cost Optimization	10 mins

Lecture and Assessment: This module provides depth in the Cost Optimization pillar of the Well-Architected Framework.

- Describe how to choose components to optimize cost
  - Make procurement recommendations to optimize costs for EC2 instances
  - Describe what happens when a spot instance is marked for termination
  - Determine the best approach for instance capacity planning and baseline utilization
  - Describe use cases for reserved instances compared with dedicated instances
  - Explain the business case and advantages for Auto Scaling
  - Identify the checks available in AWS Trusted Advisor
  - Describe the Simple monthly calculator and what it is used for
  - Describe how ElastiCache affects the performance and cost of frequently accessed data in RDS
  - Identify uses for and limits of the AWS free tier
- Describe how to provide reliability by using these AWS tools and services: Auto Scaling, Reserved Instances, Cost Allocation Tags, SNS, CloudWatch, AWS Trusted Advisor
- Analyze system resources for inefficient costs or budget overruns
  - Evaluate use cases and recommend solutions for greater efficiency and cost optimization

Part 4: Deployment and Implementation

Unit 3.12	Lecture	Troubleshooting	20 mins
Unit 3.12	**Knowledge Assessment**	Troubleshooting	10 mins

Lecture and Assessment: This module explores common scenarios reported in cloud environments including five scenarios with multiple options for identifying the root cause of the issue.

- Recommend items to check when an instance connection times out
- Recommend considerations that may affect network performance
- Recommend considerations that may affect CPU load on RDS instances
- Recommend items to check in response to an “access denied” alert
- Compare IOPS of EBS volume types

Unit 3 : Architecting on AWS (continued)

Part 4: Deployment and Implementation (continued)

Unit 3.14	Lecture	Design Patterns and Sample Architectures	30 mins
14	Lab	Using AWS Services to Enhance a Web Application	180 mins
Unit 3.14	**Knowledge Assessment**	Design Patterns and Sample Architectures	20 mins

Lecture and Assessment: This module reviews common architecture patterns and design patterns with an emphasis on typical customer-tested solutions.

- Describe a common multi-AZ pattern
- Describe a common High-Availability Database pattern
- Describe a common floating IP pattern
- Describe a common floating interface pattern
- Describe a common state-sharing pattern
- Describe a common scheduled scale-out pattern
- Describe a common job observer pattern
- Describe the purpose for using bootstrap scripts when launching an instance
- Describe use cases that leverage Elastic MapReduce
- Describe big data use cases that leverage Amazon EMR
- Describe real-time data processing use cases that leverage Amazon Kinesis
- Describe a best practice to "avoid single points of failure"
- Describe a best practice to "build security in every layer"

Lab 13: By the end of this lab, you will be able to:
- How a web application can use Active Directory as its authentication and authorization service.
- Bootstrapping an instance using S3 as its code repository.
- Connecting a web application to a RDS database and inserting and retrieving data to/from it
- Adding elasticity by using Elastic Load Balancing (ELB).
- Using Amazon CloudFront and Amazon Simple Storage Service (S3) to serve static content to a web application.

2	**Project**	Final Project	420 mins

After completing this project, you will be able to:

- Determine the region, VPCs, subnets, and Availability Zone requirements.
- Document encryption and security details.
- Design a plan for storage and backups.
- Determine how to resolve the issues concerning the Web, App, and Database Tiers.
- Build the infrastructure based on this document.

Curriculum Total **60 hours - 3600 mins**