aws training and certification

UK G-Cloud 14 Framework Agreement Amazon Web Services EMEA SARL, UK Branch (AWS) – Training Services Definition Document

May 2024

Points of Contact:

John Davies
Director
AWS UK Public Sector

Scott Powell
Business Development Manager
AWS Training and Certification

aws-gcloud@amazon.com

Table of Contents

aws training and certification

How to Use the AWS Service Definition Documents

To make it easier for customers to review AWS service content from the hundreds of individual AWS listings on the Digital Marketplace, AWS has grouped the descriptions from its listed services into bundled Service Definition Documents that describe the features of each family of AWS Cloud services. The AWS service families are:

- Cloud Compute Infrastructure Services (Lot 1 & 2)
- Professional Services (Lot 3)
- Support Services (Lot 3)
- Training Services (Lot 3)
- AWS Managed Services (Lot 3)

This AWS Training Service Definition document describes the training courses available to Customers on G-Cloud 14. Each training course has its own unique Digital Marketplace Service ID listing.

AWS has combined its training course descriptions into a consolidated document for ease of review by Customers. To access the different training courses through a Call-Off Contract, the Customer must reference each individual Digital Marketplace Service ID within the Call-Off Contract in order to enable that training course as an option that can be procured under their G-Cloud 14 Call-Off Contract. AWS would recommend that Customers list all AWS Training Services Digital Marketplace Service ID's in its Call-Off Contract to enable the option to access any training course flexibly during the term of the contract.

Once the G-Cloud 14 Call-Off Contract is executed with AWS, the parties would define the specific training course scope requirements (e.g. type and location of course, attendee numbers etc.) in one or more subsequent Statements of Work.

For a list of all AWS Digital Marketplace Service ID's, please contact an AWS account representative through aws-gcloud@amazon.com.

aws training and certification

AWS Training Offerings – High-Level Overview

Take in-person and virtual training from instructors who teach your team in-demand cloud skills in a hands-on learning environment.

Classroom training gives your team hands-on experience with new cloud skills, using a mix of presentations, peer discussion, and interactive labs. We also offer hundreds of digital training courses to help your team build new cloud skills and learn about the latest services when and where it's convenient.

SKILL DOMAIN	COURSE TITLE	DESCRIPTION	DURATION
AWS Fundamentals	AWS Technical Essentials	Learn fundamentals to become more proficient in identifying AWS services so that you can make informed decisions about IT solutions based on your business requirements and get started working on AWS.	1 day
AWS Fundamentals	AWS Cloud Practitioner Essentials	The fundamental-level full day course is intended for individuals who seek an overall understanding of the AWS Cloud, independent of specific technical roles. It provides a detailed overview of cloud concepts, AWS services, security, architecture, pricing, and support.	1 day
AWS Fundamentals	AWS Cloud Essentials for Business Leaders	Learn the fundamental concepts of cloud computing and how a cloud strategy can help companies meet business objectives. It explores the advantages and possibilities of cloud computing. It also introduces addresses concepts such as security and compliance to help facilitate better discussions with line of business (LOB) professionals and executives.	0.5 day
Architecting	Architecting on AWS	Learn the fundamentals of building IT infrastructure on AWS, and learn how to optimize the AWS Cloud by understanding AWS services and how they fit into cloud-based solutions.	3 days
Architecting	Architecting on AWS (Enhanced)	As per the 3-day course above, with additional hands-on content, and delivered at a slower pace.	4 days
Architecting	AWS Well-Architected Best Practices	The Well-Architected Framework enables you to make informed decisions about your architecture in a cloud native way and understand the impact of design decisions that are made. This course is designed to provide a deep dive into the AWS Well-Architected Framework and its 5 pillars, it covers the Well-Architected Review Process using the AWS Well-Architected Tool to complete reviews.	1 day
Architecting	Designing and Implementing Storage on AWS	AWS offers a broad portfolio of storage services and solutions with diverse capabilities for storing, accessing, and protecting your data. In this course, you will learn where, how, and when to take advantage of these different service offerings.	3 days
Architecting	Advanced Architecting on AWS	Building on Architecting on AWS, learn how to build complex solutions which incorporate data services, governance, and security on AWS.	3 days
Architecting	AWS Advanced Well-Architected Best Practices	The course covers the phases of a review, including how to prepare, run, and get guidance after a review has been performed. Attendees should have familiarity with the AWS concepts, terminology, services, and tools that are covered in the intermediate, 200-level AWS Well-Architected Best Practices	1 day

Containers	Running Containers on Amazon Elastic Kubernetes Services (EKS)	Learn container management and orchestration for Kubernetes using Amazon EKS.	3 days
Cost Management	AWS Cloud Financial Management for Builders	In this course you will learn how to implement architectural best practices, explore cost optimization strategies, and design patterns to help you architect cost-efficient solutions on AWS.	3 days
Cost Management	AWS Cloud for Finance Professionals	Gain the foundational knowledge that you need to manage, optimize and plan your cloud spend, and influence your organization's builders to be more accountable and cost-conscious.	2 days
Data Analytics	Building Batch Data Analytic Solutions on AWS	Learn to build batch data analytics solutions using Amazon EMR, an enterprise-grade Apache Spark and Apache Hadoop managed service. You will learn how Amazon EMR integrates with open-source projects such as Hive, Hue, and HBase, and with AWS services such as AWS Glue and AWS Lake Formation.	1 day
Data Analytics	Building Data Analytic Solutions using Amazon Redshift	Learn to integrate Amazon Redshift with a data lake to support both analytics and machine learning workloads. You will also learn to apply security, performance, and cost management best practices to the operation of Amazon Redshift	1 day
Data Analytics	Building Data Lakes on AWS	Learn how to build an operational data lake that supports analysis of both structured and unstructured data. You will learn the components and functionality of the services involved in creating a data lake.	1 day
Data Analytics	Data Warehousing on AWS	Learn concepts, strategies, and best practices for designing a cloud-based data warehousing solution using Amazon Redshift.	3 days
Data Analytics	Building Streaming Data Analytics Solutions on AWS	You will learn to build streaming data analytics solutions using AWS services, including Amazon Kinesis and Amazon Managed Streaming for Apache Kafka (Amazon MSK). Amazon Kinesis is a massively scalable and durable real-time data streaming service. Amazon MSK offers a secure, fully managed, and highly available Apache Kafka service.	1 day
Data Analytics	Authoring Visual Analytics Using Amazon QuickSight	QuickSight allows everyone in your organization to understand your data by exploring through interactive dashboards, asking questions in natural language, or automatically looking for patterns and outliers powered by machine learning. This course focuses on connecting to data sources, building visuals, designing interactivity, and creating calculations. You will learn how to apply security best practices to your analyses. You will also explore the machine learning capabilities built into QuickSight.	2 days
Databases	Planning and Designing Databases on AWS	Explore the key features of AWS database services and learn how to choose the appropriate AWS database service to meet your application's needs and requirements.	3 days
Databases	Building Modern Applications with AWS NoSQL Databases	In this course, you use AWS purpose-built databases to build a typical modern application with diverse access patterns and real-time scaling needs.	1 day
Developer	Developing on AWS	Learn how to use the AWS SDK to develop secure and scalable cloud applications.	3 days
Developer	Advanced Developing on AWS	Learn how to take a legacy, on premise monolithic application and refactoring it into a serverless microservices architecture.	3 days
DevOps	DevOps Engineering on AWS	Learn the most common DevOps patterns to develop, deploy, and maintain applications on the AWS platform.	3 days
Machine Learning	MLOps Engineering on AWS	This course stresses the importance of data, model and code to successful ML deployments. It will demonstrate the use of tools, automation, processes and teamwork in addressing the challenges	3 days

		associated with handoffs between data engineers, data scientists, software developers and operations.	
Machine Learning	Practical Data Science with Amazon SageMaker	Learn how to solve a real-world use case with Machine Learning (ML) and produce actionable results using Amazon SageMaker.	1 day
Machine Learning	Amazon SageMaker Studio for Data Scientists	Amazon SageMaker Studio helps data scientists prepare, build, train, deploy, and monitor machine learning (ML) models quickly by bringing together a broad set of capabilities purpose-built for ML. This course prepares experienced data scientists to use the tools that are part of SageMaker Studio to improve productivity at every step of the ML lifecycle	3 days
Machine Learning	Developing Generative AI Applications on AWS	This course is designed to introduce generative AI to software developers interested in using large language models without fine-tuning. The course provides an overview of generative AI, planning a generative AI project, getting started with Amazon Bedrock, the foundations of prompt engineering, and the architecture patterns to build generative AI applications using Amazon Bedrock and LangChain.	2 days
Migrate and Transfer	AWS Migration Essentials	This course is intended to provide solution architects with the foundational knowledge required to successfully plan and perform lift and shift migrations to the AWS Cloud. In this course you will learn about methodologies for discovering, performing and tracking migrations using various AWS tools and services.	1 day
Migrate and Transfer	Migrating to AWS	You learn about various cloud migration strategies and how to apply each step of the migration process, including portfolio discovery, application migration planning and design, performing a migration, and post-migration validation and application optimization.	3 days
Networking and Content Delivery	Networking Essentials for Cloud Applications on AWS	This course provides a comprehensive understanding of networking concepts and services within the Amazon Web Services (AWS) cloud environment. Its purpose is to equip learners with the knowledge and skills that are required to design, configure, and optimize network infrastructure on AWS.	1 day
Operations	Cloud Operations on AWS	Learn how to create automatable and repeatable deployments of networks and systems on the AWS platform.	3 days
Serverless	Developing Serverless Solutions on AWS	Learn AWS Serverless Frameworks and how to build, secure, deploy, and manage modern serverless applications.	3 days
Security	AWS Security Essentials	Learn about fundamental AWS Cloud security concepts, including AWS access control, data encryption methods, and how network access to your AWS infrastructure can be secured.	1 day
Security	AWS Security Best Practices	AWS Security Best Practices provides an overview of some of the industry best practices for using AWS security and control types. This course helps you understand your responsibilities while providing valuable guidelines for how to keep your workload safe and secure.	1 day
Security	AWS Security Governance at Scale	Governance at scale is a new concept of automating cloud governance that can help organizations retire manual processes in account management, budget enforcement and security and compliance. In this course, you will learn how to facilitate developer speed and agility and incorporate preventative and detective controls. By the end of this course, you will be able to apply governance best practices.	1 day
Security	Security Engineering on AWS	Learn how to efficiently use AWS security services to stay secure in the AWS Cloud.	3 days

AWS Training Offerings – Courses in Detail

AWS Foundational Courses

The following subsections describe foundational half or one-day courses that introduce users to the benefits and core services of the AWS Cloud.

AWS Technical Essentials

Course description

AWS Technical Essentials introduces you to essential AWS services and common solutions. The course covers the fundamental AWS concepts related to compute, database, storage, networking, monitoring, and security. You will start working in AWS through hands-on course experiences. The course covers the concepts necessary to increase your understanding of AWS services, so that you can make informed decisions about solutions that meet business requirements. Throughout the course, you will gain information on how to build, compare, and apply highly available, fault tolerant, scalable, and cost-effective cloud solutions.

Level	Delivery method	Duration
Fundamental	Instructor-led training, hands-on labs, and group exercises	1 day

Course objectives

In this course, you will learn how to:

- Describe terminology and concepts related to AWS Services
- Navigate the AWS Management Console
- Articulate the key concepts of AWS security measures and AWS Identity and Access Management (IAM)
- Distinguish among several AWS compute services, including Amazon Elastic Compute Cloud (Amazon EC2), AWS Lambda, Amazon Elastic Container Service (Amazon ECS), and Amazon Elastic Kubernetes Service (EKS).
- Understand AWS database and storage offerings, including Amazon Relational Database Service (Amazon RDS), Amazon DynamoDB, and Amazon Simple Storage Service (Amazon S3)
- Explore AWS networking services
- Access and configure Amazon CloudWatch monitoring features

Intended audience

This course is intended for:

- Individuals responsible for articulating the technical benefits of AWS services
- Individuals interested in learning how to get started with using AWS
- SysOps Administrators, Solution Architects and Developers interested in using AWS services

Prerequisites

We recommend that attendees of this course have:

- IT experience
- Basic knowledge of common data center architectures and components (servers, networking, databases, applications, and so on)
- No prior cloud computing or AWS experience required.

AWS Technical Essentials course outline

Module 1: Introduction to Amazon Web Services

- Introduction to AWS Cloud
- Security in AWS Cloud
- Hosting the employee directory application in AWS
- Hands-on Lab: Introduction to AWS Identity and Access Management (IAM)

Module 2: AWS Compute

- Compute as a service in AWS
- Introduction to Amazon Elastic Compute Cloud
- Amazon EC2 Instance lifecycle
- AWS Container Service
- What is serverless?
- Introduction to AWS Lambda
- Choose the right compute service
- Hands-on Lab: Launch the Employee Directory Application on Amazon EC2

Module 3: AWS Networking

- Networking in AWS
- Introduction to Amazon Virtual Private Cloud (Amazon VPC)
- Amazon VPC Routing
- Amazon VPC Security
- Hands-on Lab: Create a VPC and Relaunch the Corporate Directory Application in Amazon EC2

Module 4: AWS Storage

- AWS Storage types
- Amazon EC2 instance storage and Amazon Elastic Block Store (Amazon EBS)
- Object Storage with Amazon S3
- Choose the right storage service
- Hands-on Lab: Create and Amazon S3 Bucket

Module 5: Databases

- Explore databases in AWS
- Amazon Relational Database Service
- Purpose-built databases
- Introduction to Amazon DynamoDB
- Choose the right AWS database service
- Hand-on Lab: Implement and manage Amazon DynamoDB

Module 6: Monitoring, Optimization, and Serverless

- Monitoring
- Optimization

- Alternate serverless employee's directory application architecture
- Hand-on Lab: Configure High Availability for your application

Module 7: Course Summary

AWS Cloud Practitioner Essentials

Course description

This fundamental-level course is intended for individuals who seek an overall understanding of the AWS Cloud, independent of specific technical roles. You will learn about AWS cloud concepts, AWS services, security, architecture, pricing, and support to build your AWS Cloud knowledge. Throughout the day there are hands-on lab exercises to reinforce some of the core concepts of the class. It also helps you prepare for the AWS Certified Cloud Practitioner exam.

Level	Modality	Duration
Fundamental	Instructor-led training and hands-on labs	1 day

Course objectives

In this course, you will learn how to:

- Summarize the working definition of AWS
- Differentiate between on-premises, hybrid-cloud and all-in cloud
- Describe the basic global infrastructure of the AWS Cloud
- Explain the six benefits of the AWS Cloud
- Describe and provide an example of the core AWS services, including compute, network, databases and storage.
- Identify an appropriate solution using AWS Cloud services with various use cases
- Describe the Well-Architected Framework
- Explain the Shared Responsibility model
- Describe the core security services within the AWS Cloud
- Describe the basics of AWS Cloud migration
- Articulate the financial benefits of the AWS Cloud for an organization's cost management
- Define the core billing, account management, and pricing models
- Explain how to use pricing tools to make cost-effective choices for AWS services

Intended audience

This course is intended for:

- Sales and Marketing
- Legal
- Business Analysts
- Project Managers
- AWS Academy students
- IT professionals

Prerequisites

We recommend that attendees of this course have:

- General IT Business knowledge
- General IT Technical knowledge

AWS Cloud Practitioner Essentials course outline

Module 1: Introduction to Amazon Web Services

- Summarize the benefits of AWS
- Describe the differences on-demand delivery and cloud deployments
- Summarize the pay-as-you-go pricing model

Module 2: Compute in the Cloud

- Describe the benefits of Amazon Elastic Compute Cloud (Amazon EC2) at a basic level
- Identify the different Amazon EC2 Instance types
- Differentiate between the various billing options for Amazon EC2
- Describe the benefits of Amazon EC2 Auto Scaling
- Summarize the benefits of Elastic Load Balancing
- Summarize the difference between Amazon Simple Notification Service (Amazon SNS) and Amazon Simple Queue Services (Amazon SQS)
- Summarize additional AWS compute options

Module 3: Global Infrastructure and Reliability

- Summarize the benefits of the AWS Global Infrastructure
- Describe the basic concepts of Availability Zones
- Describe the benefits of Amazon CloudFront and Edge Locations
- Compare different methods of provisioning AWS services

Module 4: Networking

- Describe the basic concepts of networking
- Describe the difference between public and private networking resources
- Explain a virtual private gateway using a real life scenario
- Explain a virtual private network (VPN) using a real life scenario
- Describe the benefit of AWS Direct Connect
- Describe the benefit of hybrid deployments
- Describe the layers of security used in an IT Strategy
- Describe which services are used to interact with the AWS global network

Module 5: Storage and Databases

- Summarize the basic concept of storage and databases
- Describe the benefits of Amazon Elastic Block Store (Amazon EBS)
- Describe the benefits of Amazon Simple Storage Service (Amazon S3)
- Describe the benefits of Amazon Elastic File System (Amazon EFS)
- Summarize the various storage solutions
- Describe the benefits of Amazon Relational Database Service (Amazon RDS)
- Describe the benefits of Amazon DynamoDB
- Summarize various database services

Module 6: Security

- Explain the benefits of the shared responsibility model
- Describe Multi-factor authentication (MFA)
- Differentiate between AWS Identity and Access Management (IAM) security levels
- Describe security policies at a basic level
- Explain the benefits of AWS Organizations
- Summarize the benefits of compliance with AWS
- Explain the primary AWS security services at a basic level

Module 7: Monitoring and Analytics

- Summarize approaches to monitoring your AWS environment
- Describe the benefits of Amazon CloudWatch
- Describe the benefits of AWS CloudTrail
- Describe the benefits of AWS Trusted Advisor

Module 8: Pricing and Support

- Understand AWS pricing and support models
- Describe the AWS Free Tier
- Describe the key benefits of AWS Organizations and consolidated billing
- Explain the benefits of AWS Budgets
- Explain the benefits of AWS Cost Explorer
- Explain the primary benefits of the AWS Pricing Calculator
- Distinguish between various AWS Support Plans
- Describe the benefits of AWS Marketplace

Module 9: Migration and Innovation

- Understand migration and innovation in the AWS Cloud
- Summarize the AWS Cloud Adoption Framework (AWS CAF)
- Summarize the six key factors of a cloud migration strategy
- Describe the benefits of various AWS data migration solutions, such as AWS Snowcone, AWS Snowball and AWS Snowmobile
- Summarize the broad scope of innovative solutions that AWS offers
- Summarize the five pillars of the AWS Well-Architected Framework

Module 10: AWS Certified Cloud Practitioner Basics

- Determine resources for preparing for the AWS Certified Cloud Practitioner examination
- Describe benefits of becoming AWS Certified

AWS Cloud Essentials for Business Leaders

Course description

This course, you will learn the fundamental concepts of cloud computing and how a cloud strategy can help companies meet business objectives. It explores the advantages and possibilities of cloud computing. It also introduces and addresses concepts such as security and compliance to help facilitate better discussions with line of business (LoB) professionals and executives.

Level	Duration	Format	Delivery method
Fundamental	4 hours	Instructor-led training	Classroom, virtual classroom, and private on-site training.

Course objectives

This course is you will learn to:

- Explain the role of information technology (IT) in an organization for business transformation
- Explain the customer value proposition for using the cloud across industries
- Define key characteristics of cloud computing
- Explain the cloud business model
- Identify key security practices of cloud computing
- Frame the cloud business value using the Cloud Value Framework

Intended audience

This course is intended for:

- Line of Business (LoB) owners and executives

Prerequisites

We recommend that attendees of this course have:

- No Prior IT or cloud experience required

AWS Cloud Essentials for Business Leaders course outline

Module 1: Course Introduction

Module 2: Information Technology for Business Transformation

- Role of IT in an organization for business transformation
- Brief history of IT
- Legacy approach to IT
- What drives customers to move from traditional infrastructure to the cloud

Module 3: Cloud Computing

- Define cloud computing
- Key Characteristics of cloud technology
- The cloud business model
- Key security practices within the cloud

Module 4: Business Value of the Cloud

- The customer value proposition
- Identify who is using cloud computing
- Industry trends
- Customer examples

Module 5: The Cloud Value Framework

- Introduction to the Cloud Value Framework
- Cost savings
- Staff productivity
- Operational resilience
- Business agility

Module 6: Business Value Activity

- Using a fictional customer case study, we review and apply lessons learned from the course.

AWS Technical Courses – Grow your skills

The following subsections describe three-day instructor-led courses from which you will deepen your technical skills and learn best practices for architecting, developing, and operating infrastructure and applications on AWS

Architecting on AWS

Course description

This course focuses on the fundamentals of building IT infrastructure on the AWS platform. You will learn how to optimize the AWS Cloud by understanding AWS services and how they fit into cloud-based solutions. Best practices and design patterns are covered to help you architect optimal IT solutions on the AWS Cloud. Build and explore a variety of infrastructures through guided discussions and hands-on activity.

Level	Delivery method	Duration
Intermediate	Instructor-led training, hands-on labs, and group exercises	3 days (Standard), or, 4 days (Enhanced)

Course objectives

This course is designed to teach you how to:

- Make architectural decisions based on AWS architectural principles and best practices
- Leverage AWS services to make your infrastructure scalable, reliable, and highly available
- Leverage AWS Managed Services to enable greater flexibility and resiliency in an infrastructure
- Make an infrastructure based on AWS more efficient to increase performance and reduce costs
- Use the Well-Architected Framework to improve architectures with AWS solutions

Intended audience

This course is intended for:

- Solutions Architects
- Solution Design Engineers
- Anyone who needs to understand the scope of cloud infrastructures

Prerequisites

We recommend that attendees of this course have:

- Taken the *AWS Cloud Practitioner Essentials* course
- Working knowledge of distributed systems and multi-tier architectures
- Familiarity with general networking and cloud computing concepts

Architecting on AWS course outline

Module 1: Introduction

- The Well-Architected Framework
- AWS Global Infrastructure

Module 2: The simplest architectures

- Amazon Simple Storage Service (Amazon S3)
- Amazon S3 Glacier
- Choosing AWS Regions for your architectures
- Hands-on lab: Hosting a Static Website

Module 3: Adding a compute layer

- Amazon Elastic Compute Cloud (Amazon EC2)
- Amazon Machine Images (AMIs)
- Amazon Elastic Block Storage (Amazon EBS)
- Amazon Elastic File System (Amazon EFS)
- Amazon FSx

Module 4: Adding a database layer

- Database layer considerations
- Amazon Relational Database Service (Amazon RDS)
- Amazon DynamoDB
- AWS Database Migration Service (AWS DMS)
- Hands-on lab: Deploying a Web Application on AWS

Module 5: Networking in AWS – Part 1

- Amazon Virtual Private Cloud (Amazon VPC)
- Network security in the cloud
- Hands-on lab: Creating a VPC

Module 6: Networking in AWS – Part 2

- AWS VPN connections
- AWS Direct Connect (DX)
- VPC peering
- AWS Transit Gateway
- Load balancing on AWS
- Amazon Route 53

Module 7: AWS Identity and Access Management (IAM)

- Account users and AWS IAM
- Federating users
- Amazon Cognito

- AWS Organizations

Module 8: Elasticity, high availability, and monitoring

- Amazon CloudWatch
- AWS CloudTrail
- Amazon EC2 Auto Scaling
- Scaling your databases
- Hands-on lab: Creating a highly available environment

Module 9: Automation

- AWS CloudFormation
- AWS Systems Manager
- AWS OpsWorks
- AWS Elastic Beanstalk
- Hands-on lab: Automating infrastructure deployment with AWS CloudFormation

Module 10: Caching

- Caching on AWS with Amazon CloudFront
- Session management
- Amazon DynamoDB Accelerator (DAX)
- Amazon ElastiCache

Module 11: Building decoupled architectures

- On-Premise and cloud acquisition/deprecation cycles
- Cloud cost management tools including reporting, control and tagging
- Examples and analysis of the five pillars of cost optimization

Module 12: Microservices and serverless architectures

- Amazon Elastic Container Service (Amazon ECS)
- AWS Fargate
- AWS Lambda
- Amazon API Gateway
- AWS Step Functions
- Hands-on lab: Implementing a serverless architecture with AWS Managed Services

Module 13: RTO/RPO and backup recovery setup

- Disaster planning
- Data replication
- Recovery strategies
- AWS Storage Gateway

Module 14: Optimization and review

- Best practices for optimization
- Review questions

Architecting on AWS (Enhanced)

Course description

Architecting on AWS is for solutions architects, solution-design engineers, and developers seeking an understanding of AWS architecting. In this course, you will learn to identify services and features to build resilient, secure, and highly available IT solutions on the AWS Cloud.

Architectural solutions differ depending on industry, types of applications, and business size. AWS Authorized Instructors emphasize best practices using the AWS Well-Architected Framework, and guide you through the process of designing optimal IT solutions based on real-life scenarios. The modules focus on account security, networking, compute, storage, databases, monitoring, automation, containers, serverless architecture, edge services, and backup and recovery. At the end of the course, you will practice building a solution and apply what you have learned.

Best practices and design patterns are covered to help you architect optimal IT solutions on the AWS Cloud. Build and explore a variety of infrastructures through guided discussions and hands-on activity.

Level	Delivery method	Duration
Intermediate	Instructor-led training, hands-on labs, and group exercises	4 days

Course objectives

In this course you will learn to:

- Identify AWS architecting basic practices
- Summarize the fundamentals of account security
- Identify strategies to build a secure virtual network that includes private and public subnets
- Practice building a multi-tier architecture in AWS
- Identify strategies to select the appropriate compute resources based on business use cases
- Compare and contrast AWS storage products and services based on business scenarios
- Compare and contrast AWS database services based on business needs
- Identify the role of monitoring, load balancing, and auto scaling responses based on business needs
- Identify and discuss AWS automation tools that will help you build, maintain, and evolve your infrastructure
- Discuss hybrid networking, network peering, and gateway and routing solutions to extend and secure your infrastructure
- Explore AWS container services for the rapid implementation of an infrastructure-agnostic, portable application environment
- Identify the business and security benefits of AWS serverless services based on business examples
- Discuss the ways in which AWS edge services address latency and security
- Explore AWS backup, recovery solutions, and best practices to ensure resiliency and business continuity

Intended audience

This course is intended for:

- Solutions Architects
- Solution Design Engineers

- Developers seeking an understanding of AWS architecting
- Individuals seeking the AWS Solutions Architect-Associate certification

Prerequisites

We recommend that attendees of this course have:

- Completed AWS Cloud Practitioner Essentials, or AWS Technical Essentials
- Working knowledge of distributed systems
- Familiarity with general networking concepts
- Familiarity with IP addressing
- Working knowledge of multi-tier architectures
- Familiarity with cloud computing concepts

Architecting on AWS (Enhanced) course outline

Day One

Module 1: Architecting Fundamentals Review

- AWS services
- AWS infrastructure
- AWS Well-Architected Framework
- Hands-on lab 1: Explore using the AWS API to deploy an EC2 instance

Day Two

Module 4: Compute

- Compute services
- EC2 instances
- Storage for EC2 instances
- Amazon EC2 pricing options
- AWS Lambda
- Hands-On Lab 2: Build your Amazon VPC infrastructure

Module 5: Storage

- Storage services
- Amazon S3
- Shared file systems
- Data migration tools

Module 6: Database Services

- Database services
- Amazon RDS
- Amazon DynamoDB
- Database caching
- Database migration tools
- Hands-on Lab 3: Create a database layer in your Amazon VPC infrastructure

Day Three

Module 7: Monitoring and Scaling

- Monitoring
- Alarms and events
- Load balancing
- Auto scaling
- Hands-on Lab 4: Configure high availability in your Amazon VPC

Module 8: Automation

- AWS CloudFormation
- Infrastructure management

Module 9: Containers

- Microservices
- Containers
- Container services

Module 10: Networking 2

- VPC endpoints
- VPC peering
- Hybrid networking
- AWS Transit Gateway

Module 11: Serverless

- What is serverless?
- Amazon API Gateway
- Amazon SQS
- Amazon SNS
- Amazon Kinesis
- AWS Step Functions

Day Four

Q & A review

- Hands-on Lab 5: Build a serverless architecture

Module 12: Edge Services

- Edge fundamentals
- Amazon Route 53
- Amazon CloudFront
- DDoS protection
- AWS Outposts
- Hands-On Lab 6: Configure an Amazon CloudFront distribution with an Amazon S3 origin

Module 13: Backup and Recovery

- Disaster planning
- AWS Backup
- Recovery strategies
- Group Activity: Design an Architecture

Module 14: Course summary

Continued Learning Opportunity: On-your-own Capstone lab project
Build an AWS Multi-Tier architecture. Participants review the concepts and services learned in class and build a solution based on a scenario. The lab environment provides partial solutions to promote analysis and reflection. Participants deploy a highly available architecture. The instructor is available for consultation.

AWS Well-Architected Best Practices

Course description

The Well-Architected Framework enables you to make informed decisions about your organization's architectures in a cloud-native way and understand the impact of design decisions that are made. By using the Well-Architected Framework, you will understand the risks in your architecture and ways to mitigate them. The course is designed to provide a deep dive into the AWS Well-Architected Framework and its 5 pillars. This course also covers the Well-Architected Review Process and using the AWS Well-Architected Tool to complete the reviews.

Level	Delivery method	Duration
Intermediate	Instructor-led training, Interactive demos, hands-on labs, and knowledge checks	1 day

Course objectives

In this course, you will learn to:

- Identify the Well-Architected Framework features, design principles, design pillars and common uses.
- Apply the design principles, key services and best practices for each pillar of the Well-Architected Framework.
- Use the Well-Architected Tool to conduct Well-Architected Reviews.

Intended audience

This course is intended for:

- Technical professionals involved in architecting, building and operating AWS Solutions.

Prerequisites

We recommend that attendees of this course have:

- Knowledge of core AWS Services – AWS Cloud Practitioner Essentials Course
- Knowledge of core AWS Management Interfaces – *AWS Technical Essentials Course*
- Knowledge of core AWS design and architecture – *Architecting on AWS Course*

AWS Well-Architected Best Practices course outline

Module 1: Well-Architected Introduction

- History of Well-Architected
- Goals of Well-Architected
- What is the AWS Well-Architected Framework
- The AWS Well-Architected Tool

Module 2: Design Principles

- Operational Excellence
- Lab1: Operational Excellence
- Reliability
- Lab2: Reliability
- Security
- Lab3: Security
- Performance Efficiency
- Lab4: Performance Efficiency
- Cost Optimization
- Lab5: Cost Optimization

Designing and Implementing Storage on AWS

Course description

AWS offers a broad portfolio of storage services and solutions with diverse capabilities for storing, accessing, and protecting your data. In this course, you will learn where, how, and when to take advantage of these different service offerings. You will learn which services to consider when looking to solve your data storage challenges. You will learn how to best evaluate your options in selecting the appropriate AWS storage service to meet your use case and business requirements. You will also gain a better understanding of how to store, manage, and protect your data in the cloud. Through a series of hands-on exercises that demonstrate the ease and power of AWS platform, you will learn how to quickly provision powerful storage solutions in minutes.

Level	Delivery method	Duration
Intermediate	Instructor-led training, hands-on labs, and group exercises	3 days

Course objectives

In this course you will learn to:

- Describe the benefits of the core AWS storage services and identify some of their primary use cases
- Select and design an appropriate storage solution according to application and business requirements
- Configure storage resources to work with the broad array of AWS service offerings

- Select the right method to move data between on-premises workloads and the AWS Cloud
- Design storage solutions to protect data at rest and in transit
- Set up monitoring and observability for Cloud storage to gain insight into access patterns, utilization, and efficiency
- Design and optimize storage solutions according to cost, scalability, and performance requirements

Intended audience

This course is intended for:

- Solution Architects
- Cloud Storage Engineers
- Cloud Operations Specialists
- DevOps Engineers

Prerequisites

We recommend that attendees of this course have:

- Completed AWS Cloud Practitioner Essentials, AWS Technical Essentials, or equivalent real-world experience

Designing and Implementing Storage on AWS course outline

Day One

Module 1: Introduction to Cloud Storage

- Storage in the AWS cloud
- Designing Well Architected Storage Solutions
- Designing Durable and Available Storage Solutions
- Building Accessible and Secure Storage Solutions

Module 2: Designing Object Storage Solutions in AWS

- What is object storage?
- Planning and designing your Amazon S3 deployment
- Managing Amazon S3
- Access Control with Amazon S3
- Hands-On Lab: Exploring S3 Access Control and S3 Object Lambda

Module 3: Implementing Object Storage solutions with S3

- Cost management and the data lifecycle
- Managing data transfers into Amazon S3
- Data protection in Amazon S3
- Manage objects stored in Amazon S3 at scale
- Hands-on Lab: Multi-Part Uploads, Batch Operations, and Cross-Region Replication with Amazon S3

Module 4: Designing Block Storage Solutions in AWS

- Block storage fundamentals

- Amazon Elastic Block Store (Amazon EBS)
- Configuring EBS volume types
- EC2 and EBS encryption

Day Two

Module 5: Implementing Block Storage Solutions with Amazon EBS

- Creating EBS volumes
- Managing EBS volumes
- Managing EBS snapshots at scale
- Hands-On Lab: Managing EBS Volumes: Capacity, Performance, and Data Protection

Module 6: File Storage and Amazon EFS

- Cloud-based file storage
- Amazon EFS overview
- Accessing Amazon EFS
- Securing and protecting Amazon EFS file systems
- Hands-On Lab: Using Amazon EFS with AWS Lambda and Amazon ECS

Module 7: Cloud file storage with Amazon FSx

- Amazon FSx overview
- Amazon FSx for Windows File Server
- Amazon FSx for NetAPP ONTAP
- Amazon FSx for NetAPP ONTAP
- Amazon FSx for OpenZFS
- Amazon FSx for Lustre
- Choosing an Amazon FSx service
- Hands-On Lab: Working with FSx for NetApp ONTAP and FSx for OpenZFS

Day Three

Module 8: Hybrid and Edge Cloud Storage

- Hybrid and edge cloud storage overview
- Introduction to AWS Storage Gateway
- AWS Storage Gateway architectures
- AWS Snow Family

Module 9: Moving data to AWS

- Moving data to AWS
- Working with AWS DataSync
- Implementing AWS Transfer Family
- Hands-On Lab: Moving Data with Storage Gateway and DataSync

Module 10: Backup and Disaster Recovery with AWS

- Designing a data protection strategy
- AWS Backup
- Creating backup plans

- Working with AWS DRS
- Hands-On Lab: Creating and Restoring Backups with AWS Backup

Module 11: Monitoring, Automating, and Optimizing your AWS Storage

- AWS Observability Services
- Amazon S3 Storage Lens
- Amazon CloudWatch
- AWS CloudTrail
- AWS Config
- AWS Compute Optimizer
- Hands-On Lab: Storage Monitoring, Automation, and Optimization

Advanced Architecting on AWS

Course description

In this course, each module presents a scenario with an architectural challenge to be solved. You will examine available AWS services and features as solutions to the problem. You will gain insights by participating in problem-based discussions and learning about the AWS services that you could apply to meet the challenges. Over 3 days, the course goes beyond the basics of a cloud infrastructure and covers topics to meet a variety of needs for AWS customers. Course modules focus on managing multiple AWS accounts, hybrid connectivity and devices, networking with a focus on AWS Transit Gateway connectivity, container services, automation tools for continuous integration/continuous delivery (CI/CD), security and distributed denial of service (DDoS) protection, data lakes and data stores, edge services, migration options, and managing costs. The course concludes by presenting you with scenarios and challenging you to identify the best solutions.

Level	Delivery method	Duration
Advanced	Instructor-led training, hands-on labs, and group exercises	3 days

Course objectives

This course is designed to teach you how to:

- Review the AWS Well-Architected Framework to ensure understanding of best cloud design practices by responding to poll questions while following a graphic presentation
- Demonstrate the ability to secure Amazon Simple Storage Service (Amazon S3) virtual private cloud (VPC) endpoint connections in a lab environment
- Identify how to implement centralized permissions management and reduce risk using AWS Organizations organizational units (OUs) and service control policies (SCPs) with AWS Single Sign-On
- Compare the permissions management capabilities of OUs, SCPs, and AWS SSO with and without AWS Control Tower to determine best practices based on use cases
- Discuss AWS hybrid network designs to address traffic increases and streamline remote work while ensuring FIPS140-2 Level 2, or Level 3 security compliance
- Explore the solutions and products available to design a hybrid infrastructure, including access to 5Gnetworks, to optimize service and reduce latency while maintaining high security for critical on-premises applications

- Explore ways to simplify the connection configurations between applications and high-performance workloads across global networks
- Demonstrate the ability to configure a transit gateway in a lab environment
- Identify and discuss container solutions and define container management options
- Build and test a container in a lab environment
- Examine how the AWS Developer tools optimize the CI/CD Pipeline with updates based on near-real-time data
- Identify the anomaly detection and protection services that AWS offers to defend against DDoS attacks
- Identify ways to secure data in transit, at rest, and in use with AWS Key Management Services (AWS KMS) and AWS Secrets Manager
- Determine the best data management solution based on frequency of access, and data query and analysis needs
- Set up a data lake and examine the advantages of this type of storage configuration to crawl and query data in a lab environment
- Identify solutions to optimize edge services to eliminate latency, reduce inefficiencies, and mitigate risks
- Identify the components used to automate the scaling of global applications using geolocation and traffic control
- Deploy and activate an AWS Storage Gateway file gateway and AWS DataSync in a lab environment
- Review AWS cost management tools to optimize costs while ensuring speed and performance
- Review migration tools, services, and processes that AWS provides to implement effective cloud operation models based on use cases and business needs
- Provide evidence of your ability to apply the technical knowledge and experience gained in the course to improve business practices by completing a Capstone Project

Intended audience

This course is intended for:

- Solution Architects
- Anyone who designs solutions for cloud infrastructures

Prerequisites

We recommend that attendees of this course have:

- Knowledge and experience with core AWS Services from Compute, Storage, Networking and AWS Identity and Access Management (IAM) Categories
- Attended ***Architecting on AWS*** course OR
- Achieved the AWS Certified Solutions Architect – Associate certification OR
- Have at least 1-year experience operating AWS Workloads

Advanced Architecting on AWS course outline

Day One

Module 1: Reviewing Architecting Concepts

- Group Exercise: Review ***Architecting on AWS*** Core best practices
- Lab1: Securing Amazon S3 VPC Endpoint Communications

Module 2: Single to Multiple Accounts

- AWS Organizations for multi-account access and permissions
- AWS SSO to simplify access and authentication across AWS accounts and third-party services
- AWS Control Tower
- Permissions, access, and authentication

Module 3: Hybrid Connectivity

- AWS Client VPC Authentication and control
- AWS Site-to-Site VPN
- AWS Direct Connect for hybrid public and private connections
- Increasing bandwidth and reducing cost
- Basic, high and maximum resiliency
- Amazon route 53 Resolver DNS Resolution

Module 4: Specialized Infrastructure

- AWS Storage Gateway solutions
- On-demand VMware Cloud on AWS
- Extending cloud infrastructure services with AWS Outposts
- AWS Local Zones for latency-sensitive workloads
- Your 5G network with and without AWS Wavelength

Module 5: Connecting Networks

- Simplifying private subnet connections
- VPC isolation with a shared services VPC
- Transit Gateway Network Manager and VPC Reachability Analyzer
- AWS Resource Access Manager
- AWS PrivateLink and endpoint services
- Lab 2: Configuring Transit Gateways

Day Two

Module 6: Containers

- Container solutions compared to virtual machines
- Docker benefits, components, solutions architecture, and versioning
- Container hosting on AWS to reduce cost
- Managed container service, Amazon Elastic Container Service (Amazon ECS) and Amazon Elastic Kubernetes Service (Amazon EKS)
- AWS Fargate
- Lab 3: Deploying an Application with Amazon EKS on Fargate

Module 7: Continuous Integration /Continuous Delivery (CI/CD)

- CI/CD Solutions and impact
- CI/CD automation with AWS CodePipeline
- Deployment Models
- AWS CloudFormation StackSets to improve deployment management

Module 8: High Availability and DDoS Protection

- Common DDoS attacks layers
- AWS WAF
- AWS WAF web access control lists (ACL's), real-time metrics, logs, security automation
- AWS Shield Advanced services and AWS DDoS Response Team (DRT) services
- AWS Network Firewall and AWS Firewall Manager to protect accounts at scale

Module 9: Securing Data

- What cryptography is, why you would use it and how you would use it
- AWS KMS
- AWS CloudHSM architecture
- FIPS 140-2 Level 2 and Level 3 encryption
- Secrets Manager

Module 10: Large-Scale Data Stores

- Amazon S3 data storage management including storage class, inventory, metrics, and policies
- Data lake vs. Data warehouse: Differences, benefits, and examples
- AWS Lake Formation solutions, security, and control
- Lab 4: Setting Up a Data Lake with Lake Formation

Day Three

Module 11: Large-Scale applications

- What edge services are and why you would use them
- Improve performance and mitigate risk with Amazon CloudFront
- Lambda@Edge
- AWS Global Accelerator: IP Addresses, intelligent traffic distribution and health checks
- Lab 5: Migrating an On-Premises NFS Share using AWS DataSync and Storage Gateway

Module 12: Optimizing Cost

- On-Premise and cloud acquisition/deprecation cycles
- Cloud cost management tools including reporting, control and tagging
- Examples and analysis of the five pillars of cost optimization

Module 13: Migrating Workloads

- Business drivers and the process for migration
- Successful customer practices
- The 7 Rs to migrate and modernize
- Migration tools and services from AWS
- Migrating databases and large data stores
- AWS Schema Conversion Tool (AWS SCT)

Module 14: Capstone Project

- Online Course Supplement (OCS) to review use cases, investigate data and answer architecting design questions about Transit Gateway, Hybrid Connectivity, Migration and Cost Optimization

AWS Advanced Well-Architected Best Practices

Course description

This hands-on, advanced technical-level, instructor-led course provides a deep dive into Amazon Web Services (AWS) best practices to help you perform effective and efficient AWS Well-Architected Framework reviews. The course covers the phases of a review, including how to prepare, run, and get guidance after a review has been performed. This course is designed for AWS customers and AWS Partners. Attendees should have familiarity with the AWS concepts, terminology, services, and tools that are covered in the intermediate 200-level precursor to this course. This course provides an AWS Well-Architected Framework review simulation and instructor-led group exercises and discussions about prioritizing and resolving risks. The content focuses on how to prepare proposals on high and medium risk issues using the AWS Well Architected Tool.

Level	Delivery method	Duration
Advanced	Instructor-led training, Interactive demos, hands-on labs, and knowledge checks	1 day

Course objectives

In this course, you will learn to:

- Recognize Workload definition and key concepts
- Identify the AWS Well-Architected Framework review phases, process, best practices and anti-patterns
- Identify high and medium risks
- Prioritize improvements to the AWS Well-Architected workflow
- Locate and use AWS Well-Architected Framework white paper, labs and prebuilt solutions in the AWS solutions library
- Locate and use AWS Well-Architected independent software vendors (ISVs)
- Locate and use the Well-Architected Partner Program (WAPP)

Intended audience

This course is intended for:

- Technical professionals involved in architecting, building and operating AWS Solutions.

Prerequisites

We recommend that attendees of this course have:

- Completed AWS Well-Architected Best Practices

AWS Advanced Well-Architected Best Practices course outline

Module 1: AWS Well-Architected Framework Reviews

- AWS Well-Architected Framework workload
- AWS Well-Architected Framework review phases
- AWS Well-Architected review approach, lessons learned, and use cases
- AWS Well-Architected review best practices
- AWS Well-Architected review anti-patterns
- Knowledge check

Module 2: Customer Scenario Group Sessions

- Customer Story
- Demonstration of the workflow
- Hands-on Group Exercise
- Demonstration: Running a review in the Operational Excellence pillar
- Role-play exercise: Running a review in the Security pillar
- Role-play exercise: Running a review in the Reliability pillar
- Role-play exercise: Running a review in the Performance Efficiency pillar
- Role-play exercise: Running a review in the Cost Optimization pillar

Module 3: Risk Solutions and Priorities

- AWS Well-Architected Framework review engagement workflow
- High risk and medium risk issues
- Defining risks
- Resolving high-risk issues (HRIs) and medium-risk issues (MRIs)
- Group discussion: Identifying and resolving significant risks for:
    - Operational Excellence
    - Security
    - Reliability
    - Performance Efficiency
    - Cost Optimization
    - Prioritizing improvements
- AWS Well-Architected improvement workflow

Module 4: Resources

- Resource pages
- AWS Well-Architected ISVs
- AWS Well-Architected Partner Programs (WAPP)

Module 5: Course Summary

- Debrief
- What's next?
- Course feedback

Running Containers on Amazon Elastic Kubernetes Service (Amazon EKS)

Course description

In this course, you will learn how to use Amazon EKS to manage and orchestrate containers with Kubernetes. With Amazon EKS you can run Kubernetes on AWS without needing to install, operate, and maintain your own Kubernetes control plane. You will manage container images using Amazon Elastic Container Registry (Amazon ECR) and learn how to automate application deployment. You will deploy applications using continuous integration and delivery (CI/CD) tools. You will learn how to monitor and scale your environment by using metrics, logging, tracing, and horizontal and vertical scaling. You will also manage storage for your containerized applications, configure AWS networking services to support the cluster, and learn how to secure your Amazon EKS environment

Level	Delivery method	Duration
Intermediate	Instructor-led training, hands-on labs, and group exercises	3 days

Course objectives

- Describe the main components of Kubernetes, including the key objects and the core components of the Kubernetes API.
- Describe how Amazon EKS manages the Kubernetes control plane and parts of the data plane.
- Build and maintain an Amazon EKS cluster.
- Deploy applications to an Amazon EKS cluster.
- Manage applications running in enterprise-scale Amazon EKS clusters.
- Configure efficient, secure communication both within the cluster and with outside services.
- Configure observability in an Amazon EKS cluster.
- Provision storage for applications running on Amazon EKS.
- Secure an Amazon EKS cluster.
- Upgrade your Kubernetes, Amazon EKS and third-party tools

Intended audience

This course is intended for:

- Those who will provide container orchestration management in the AWS cloud including:
- Cloud Architects
- System Administrators
- DevOps Engineers

Prerequisites

We recommend that attendees of this course have:

- Completed *Introduction to Containers*
- Completed *Amazon Elastic Kubernetes Service (EKS) Primer*
- Completed *AWS Cloud Practitioner Essentials* (or equivalent real-world experience)
- Basic Linux administration experience
- Basic network administration experience

Running Containers on Amazon Elastic Kubernetes Service (Amazon EKS) course outline

Day One

Module 1: Kubernetes Fundamentals

- Benefits of containers
- Container orchestration
- Kubernetes internals
- Pod scheduling
- Kubernetes objects

Module 2: Amazon EKS Fundamentals

- Introduction to Amazon EKS
- Amazon EKS control plane
- Amazon EKS data plane
- Fundamentals of Amazon EKS security
- Two APIs: Kubernetes and Amazon EKS
- Hands-On Lab: Deploying Kubernetes Pods

Module 3: Building and maintaining an Amazon EKS cluster

- Creating an Amazon EKS cluster
- Deploying nodes
- Planning for an upgrade
- Upgrading your Kubernetes version

Module 4: Deploying Applications to Your Amazon EKS Cluster

- Application deployment methods
- Working with Amazon ECR
- Deploying applications with Helm
- Hands-On Lab: Deploying Applications

Day Two

Module 5: Managing Applications at Scale in Amazon EKS

- Scale to meet demand in Amazon EKS
- Continuous deployment in Amazon EKS
- GitOps and Amazon EKS
- Hands-On Lab: Continuous Deployment and GitOps

Module 6: Managing Networking in Amazon EKS

- Review: Networking in AWS
- Communicating in Amazon EKS
- Improving Pod-level security
- Load balancing with Services

Module 7: Configuring Observability in Amazon EKS

- Configuring observability in an Amazon EKS cluster
- Collecting metrics
- Managing logs
- Application tracing in Amazon EKS
- Hands-On Lab: Monitoring Amazon EKS

Day Three

Module 8: Managing Storage in Amazon EKS

- Design patterns for storage
- Persistent storage in Kubernetes
- Persistent storage with AWS storage services
- Managing secrets
- Hands-On Lab: Persistent Storage in Amazon EKS

Module 9: Managing Security in Amazon EKS

- Cloud security fundamentals
- Authentication and authorization
- Managing IAM and RBAC
- Managing Pod permissions using RBAC service accounts
- Hands-On Lab: Capstone Exercise

AWS Cloud Financial Management for Builders

Course description

This course is for individuals who seek an understanding of how to manage, optimize and predict costs as you run workloads on AWS. In this course you will learn how to implement architectural best practices, explore cost optimization strategies, and design patterns to help you architect cost-efficient solutions on AWS.

Level	Delivery method	Duration
Intermediate	Instructor-led training, hands-on labs, and group exercises	3 days

Course objectives

In this course you will learn to:

- Explain the cost of core AWS services
- Determine and predict costs associates with current and future cloud workloads
- Use strategies and best practices to reduce AWS costs
- Use AWS tools to manage, monitor, alert, and optimize your AWS spend
- Apply strategies to monitor service costs and usage
- Implement governance standards, including resource tagging, account structure provisioning, permissions, and access

Intended audience

This course is intended for:

- Solutions Architects
- Developers
- Cost-optimization leads
- Systems administrators

Prerequisites

We recommend that attendees of this course have:

- Taken the *Architecting on AWS* course

AWS Cloud Financial Management for Builders course outline

Day One

Module 1: Introduction to Cloud Financial Management

- Introduction to Cloud Financial Management Overview
- Four pillars of Cloud Financial Management

Module 2: Resource Tagging

- Tagging Resources
- Hands-on Lab: Cost Optimization – Control resource consumption using Tagging and AWS Config

Module 3: Pricing and Cost

- Fundamentals of pricing
- AWS Free Tier
- Volume discounts
- Savings plans and Reserved Instances
- Demonstration: AWS Pricing Calculator

Module 4: Billing, Reporting, and Monitoring

- Understanding AWS invoices
- Reporting and planning
- AWS Cost Explorer
- AWS Budgets
- Demonstration: AWS Billing Console
- Demonstration: AWS Cost Explorer
- Demonstration: Trusted Advisor
- Hands-On Lab: Deploy Ephemeral Environments using Amazon EC2 AutoScaling

Module 5: Architecting for Cost: Compute

- Evolution of compute efficiency
- Amazon EC2 right-sizing
- Purchasing options
- Architect for Amazon EC2 Spot Instance
- Impact of software licensing
- Demonstration: Compute Optimizer
- Demonstration: Spot Instance Advisor
- Hands-On Lab: Cost Optimization: EC2 Instance right-sizing using AWS CloudWatch metrics

Module 6: Architecting for Cost: Networking

- Data transfer costs

- Understand data costs for different services
- How to triage network costs
- Hands-On Lab: Cost optimization: Reduce Data Transfer Costs Using Amazon CloudFront and Endpoints

Module 7: Architecting for Cost: Storage

- Amazon EBS cost, pricing, and best practices
- Amazon S3 cost, pricing, and best practices
- Amazon EFS cost, pricing, and best practices
- Hands-On Lab: Cost optimization: Reduce Storage Costs Using Amazon S3 Lifecycle Management

Module 8: Architecting for Cost: Databases

- Amazon RDS cost, pricing, and best practices
- Amazon Aurora cost, pricing, and best practices
- Amazon DynamoDB cost, pricing, and best practices
- Amazon ElastiCache cost, pricing, and best practices
- Amazon Redshift cost, pricing, and best practices

Module 9: Cost Governance

- Setting up AWS Organizations
- AWS Systems Manager
- Hands-On Lab: Cost optimization: Reduce Compute Costs Using AWS Instance Scheduler

Module 10: Course Summary

AWS Cloud for Finance Professionals

Course description

In this course, you learn how finance professionals can use Amazon Web Services (AWS) to move to the cloud in a fiscally responsible manner. You gain foundational knowledge to help you manage, optimize, and plan your cloud spend. Lastly, you learn how to influence your organization's builders to be more accountable and cost conscious.

Level	Delivery method	Duration
Fundamental	Instructor-led training, hands-on labs, and group exercises	2 days

Course objectives

This course is you will learn to:

- Define cloud business value
- Estimate costs associated with current and future cloud workloads
- Use tools to report, monitor, allocate, optimize and plan AWS Spend
- Optimize cloud spending and usage through pricing models
- Establish best practices with Cloud Financial Management (CFM) and Cloud Financial Operations (Cloud FinOps)

- Implement financial governance and controls
- Drive finance organization innovation

Intended audience

This course is intended for enterprise finance stakeholders who want to learn how to maximize cloud business value, use CFM best practices, and help finance teams innovate with AWS.

Prerequisites

We recommend that attendees of this course have completed:

- Cloud Computing and AWS from the digital version of AWS Cloud for Finance Professionals
- AWS Cloud Practitioner Essentials course

AWS Cloud for Finance Professionals course outline

Day One

Module 1: Introduction

- Cloud Spending Decisions
- AWS Pricing
- Cost Drivers
- AWS Well-Architected Framework
- Activity 1.1: Cloud Value Metrics
- Cloud Financial Management
- Activity 1.2: Cloud Financial Management outcomes

Module 2: Planning and Forecasting

- Estimate cloud workload costs
- Activity 2.1: Build and refine a cost estimate
- Budget and forecast cloud costs
- Improve financial predictability

Module 3: Measurement and Accountability

- KPI's and unit metrics
- Cost visibility and monitoring
- Demonstration 3.1: Tools for cost visibility, tools for cost monitoring
- Cost allocation and accountability
- Cost allocation building blocks

Day Two

Module 4: Cost Optimization

- Usage optimizations
- Commitment-based purchase options
- Activity 4.1: Cost optimization

Module 5: Cloud Financial Operations

- Organizational change for Cloud Financial Management
- Organizational models for Cloud Financial Management
- Activity 5.1: Organizational models
- Establishing a cost-aware organizational culture

- Governance, control and agility
- AWS governance and control building blocks
- Automated-based governance using AWS services

Module 6: Financial Transformation and Innovation

- Keys to financial innovation
- Financial transformation
- Activity 6.1: Solutions for financial innovation

Module 7: Resources and Next Steps

Building Batch Data Analytics Solutions on AWS

Course description

In this course, you will learn to build batch data analytics solutions using Amazon EMR, an enterprise-grade Apache Spark and Apache Hadoop managed service. You will learn how Amazon EMR integrates with open-source projects such as Apache Hive, Hue, and HBase, and with AWS services such as AWS Glue and AWS Lake Formation. The course addresses data collection, ingestion, cataloging, storage, and processing components in the context of Spark and Hadoop. You will learn to use EMR Notebooks to support both analytics and machine learning workloads. You will also learn to apply security, performance, and cost management best practices to the operation of Amazon EMR.

Level	Delivery method	Duration
Intermediate	Instructor-led training, presentations, demonstrations and assessments	1 day

Course objectives

In this course, you will learn to:

- Compare the features and benefits of data warehousing, data lakes and modern data architectures
- Design and implement a batch data analytics solution
- Identify and apply appropriate techniques, including compression, to optimize data storage
- Select and deploy appropriate options to ingest, transform and store data
- Choose the appropriate instance and node types, clusters, auto scaling and network topology for a particular business use case
- Understand how data storage and processing affect the analysis and visualization mechanisms needed to gain actionable business insights
- Secure data at rest and in transit
- Monitor analytics workloads to identify and remediate problems
- Apply cost management best practices

Intended audience

This course is intended for:

- Data Platform Engineers
- Architects and operators who build and manage data analytics pipelines

Prerequisites

Students with a minimum of 1 years' experience managing open-source data frameworks, such as Apache Spark or Apache Hadoop will benefit from this course.

We recommend that attendees of this course have completed:

- Either AWS Technical Essentials or Architecting on AWS
- Either Building Data Lakes on AWS or Getting Started with AWS Glue

Building Batch Data Analytics Solutions on AWS Course Outline

Module A: Overview of Data Analytics and the Data Pipeline

- Data Analytics use cases
- Using the Data Pipeline for Data Analytics

Module 1: Introduction to Amazon EMR

- Using Amazon EMR in analytics solutions
- Amazon EMR cluster architecture
- Interactive Demo 1: Launching an Amazon EMR Cluster
- Cost management Strategies

Module 2: Data Analytics Pipeline Using Amazon EMR: Ingestion and Storage

- Storage Optimization with Amazon EMR
- Data Ingestion Techniques

Module 3: High-Performance Batch Data Analytics Using Apache Spark on Amazon EMR

- Apache Spark on Amazon EMR use cases
- Why Apache Spark on Amazon EMR
- Spark Concepts
- Interactive Demo 2: Connect to an EMR cluster and perform Scala commands using the Spark Shell
- Transformation, processing and analytics
- Using notebooks with Amazon EMR
- Practice Lab 1: Low-latency data analytics using Apache Spark on Amazon EMR

Module 4: Processing and Analyzing Batch Data with Amazon EMR and Apache Hive

- Using Amazon EMR with Hive to process batch data
- Transformation, processing and analytics
- Practice Lab 2: Batch data processing using Amazon EMR with Hive
- Introduction to Apache HBase on Amazon EMR

Module 5: Serverless Data Processing

- Serverless data processing, transformation and analytics
- Using AWS Glue with Amazon EMR workloads
- Practice Lab 3: Orchestrate data processing in Spark using AWS Step Functions

Module 6: Security and Monitoring of Amazon EMR Clusters

- Securing EMR Clusters
- Interactive Demo 3: Client-side encryption with EMRFS
- Monitoring and troubleshooting Amazon EMR Clusters
- Demo: Reviewing Apache Spark Cluster history

Module 7: Designing Batch Data Analytics Solutions

- Batch Data analytics use cases
- Activity: Designing a batch data analytics workflow

Module B: Developing Modern Data Architectures on AWS

- Modern Data Architectures

Building Data Analytics Solutions Using Amazon Redshift

Course description

In this course you will build a data analytics solution using Amazon Redshift, a cloud data warehouse service. The course focuses on the data collection, ingestion, cataloging, storage and processing components of the analytics pipeline. You will learn to integrate Amazon Redshift with a data lake to support both analytics and machine learning workloads. You will also learn to apply security, performance and cost management best practices to the operation of Amazon Redshift.

Level	Delivery method	Duration
Intermediate	Instructor-led training, Interactive demos, hands-on labs, and group exercises.	1 day

Course objectives

In this course, you will learn to:

- Compare the features and benefits of data warehouse, data lakes and modern data architecture
- Design and implement a data warehouse analytics solution
- Identify and apply appropriate options to ingest, transform and store data
- Choose the appropriate instance and node types, clusters, auto scaling and network topology for a particular business use case
- Understand how data storage and processing affect the analysis and visualization mechanisms needed to gain actionable business insights
- Secure data at rest and in transit
- Monitor analytics workloads to identify and remediate problems
- Apply cost management best practices

Intended audience

This course is intended for data warehouse engineers, data platform engineers, and architects and operators who build and manage data analytics pipelines

Prerequisites

We recommend that attendees of this course have:

- Completed the AWS Technical Essentials or Architecting on AWS course
- Completed Building Data Lakes on AWS course

Building Data Analytics Solutions Using Amazon Redshift Course Outline

Module A: Overview of Data Analytics and Data Pipeline

- Data Analytic use cases
- Using the data pipeline for analytics

Module 1: Using Amazon Redshift in the Data Analytics Pipeline

- Why Amazon Redshift for Data Warehousing?
- Overview of Amazon Redshift

Module 2: Introduction to Amazon Redshift

- Amazon Redshift architecture
- Interactive Demo 1: Touring the Amazon Redshift console
- Amazon Redshift features
- Practice Lab 1: Load and query data in an Amazon Redshift cluster

Module 3: Ingestion and Storage

- Ingestion
- Interactive Demo 2: Connecting your Amazon Redshift cluster using a Jupyter notebook with Data API
- Data distribution and storage
- Interactive Demo 3: Analyzing semi-structured date using SUPER data type
- Querying data in Amazon Redshift
- Practice Lab 2: Data Analytics using Amazon Redshift Spectrum

Module 4: Processing and Optimizing Data

- Data Transformation
- Advanced Querying
- Practice Lab 3: Data transformation and querying in Amazon Redshift
- Resource Management
- Interactive Demo 4: Applying mixed workload management on Amazon Redshift
- Automation and optimization
- Interactive Demo 5: Amazon Redshift Cluster resizing from the dc2.large to ra3.xplus cluster

Module 5: Security and Monitoring of Amazon Redshift Clusters

- Securing the Amazon Redshift Cluster

- Monitoring and troubleshooting Amazon Redshift Clusters

Module 6: Designing Data Warehouse Analytics Solutions

- Data Warehouse use case review
- Activity: Designing a data warehouse analytics workflow

Module B: Developing Modern Data Architectures on AWS

- Modern data architectures

Building Data Lakes on AWS

Course description

In this course you will learn how to build an operational data lake that supports analysis of both structured and unstructured data. You will learn the components of functionality of the services involved in creating a data lake. You will use AWS Lake Formation to build a data lake. AWS Glue to build a data catalog, and Amazon Athena to analyze data. The course lectures and labs further your learning with the exploration of several common data lake architectures.

Level	Delivery method	Duration
Intermediate	Instructor-led training, hands-on labs, and group exercises.	1 day

Course objectives

In this course, you will learn to:

- Apply data lake methodologies in planning and designing a data lake
- Articulate the components and services required for building an AWS Data Lake
- Secure a data lake with appropriate permissions
- Ingest, store and transform data in a data lake
- Query, Analyze and visualize data within a data lake

Intended audience

This course is intended for:

- Data platform engineers
- Solutions Architecture
- IT professionals

Prerequisites

We recommend that attendees of this course have:

- Completed the AWS Technical Essentials course
- One year of experience building data analytics pipelines or have completed the *Data Analytics Fundamentals* digital course

Building Data Lakes on AWS Course Outline

Module 1: Introduction to Data Lakes

- Describe the value of data lakes
- Compare Data lakes and data warehouses
- Describe the components of a data lake
- Recognize common architectures built on data lakes

Module 2: Data ingestion, cataloging and preparation

- Describe the relationship between data lake storage and data ingestion
- Describe AWS Glue Crawlers and how they are used to create a data catalog
- Identify data formatting, partitioning and compression for efficient storage and query
- Lab 1: Set up a simple data lake

Module 3: Data processing and analytics

- Recognize how data processing applies to a data lake
- Use AWS Glue to process data within a data lake
- Describe how to use Amazon Athena to analyze data in a data lake

Module 4: Building a data lake with AWS Lake Formation

- Describe the features and benefits of AWS Lake Formation
- Use AWS Lake Formation to create a data lake
- Understand the AWS Lake Formation Security Model
- Lab 2: Build a data lake using AWS Lake Formation

Module 5: Additional Lake Formation configurations

- Automate AWS Lake Formation using blueprints and workflows
- Apply security and access controls to AWS Lake Formation
- Match records with AWS Lake Formation FindMatches
- Visualize data with Amazon QuickSight
- Lab 3: Automate data lake creation using AWS Lake Formation blueprints
- Lab 4: Data visualization using Amazon QuickSight

Module 6: Architecture and Course Review

- Post course knowledge check
- Architecture review
- Course review

Data Warehousing on AWS

Course description

Data Warehousing on AWS introduces you to concepts, strategies, and best practices for designing a cloud-based data warehousing solution using Amazon Redshift, the petabyte-scale data warehouse in AWS. This course demonstrates how to collect, store, and prepare data for the data warehouse by using AWS services such as Amazon DynamoDB, Amazon EMR, Amazon Kinesis, and Amazon S3. Additionally, this course demonstrates how to use Amazon QuickSight to perform analysis on your data.

Level	Delivery method	Duration
Intermediate	Instructor-led training, hands-on labs and group exercises	3 days

Course objectives

This course is designed to teach you how to:

- Discuss the core concepts of data warehousing, and the intersection between data warehousing and big data solutions
- Launch an Amazon Redshift cluster and use the components, features, and functionality to implement a data warehouse in the cloud
- Use other AWS data and analytic services, such as Amazon DynamoDB, Amazon EMR, Amazon Kinesis, and Amazon S3, to contribute to the data warehousing solution
- Architect the data warehouse
- Identify performance issues, optimize queries, and tune the database for better performance
- Use Amazon Redshift Spectrum to analyze data directly from an Amazon S3 bucket
- Use Amazon QuickSight to perform data analysis and visualization tasks against the data warehouse

Intended audience

This course is intended for:

- Database Architects
- Database Administrators
- Database Developers
- Data Analysts
- Data Scientists

Prerequisites

We recommend that attendees of this course have:

- Taken *AWS Technical Essentials* Course (Or Equivalent AWS Experience)
- Familiarity with relational databases and database design concepts

Data Warehousing on AWS course outline

Day One

Module 1: Introduction to Data Warehousing

- Relational databases
- Data warehousing concepts
- The intersection of data warehousing and big data
- Overview of data management in AWS
- Hands-on lab 1: Introduction to Amazon Redshift

Module 2: Introduction to Amazon Redshift

- Conceptual overview
- Real-world use cases
- Hands-on lab 2: Launching an Amazon Redshift cluster

Module 3: Launching clusters

- Building the cluster
- Connecting to the cluster
- Controlling access
- Database security
- Load data
- Hands-on lab 3: Optimizing database schemas

Day Two

Module 4: Designing the database schema

- Schemas and data types
- Columnar compression
- Data distribution styles
- Data sorting methods

Module 5: Identifying data sources

- Data sources overview
- Amazon S3
- Amazon DynamoDB
- Amazon EMR
- Amazon Kinesis Data Firehose
- AWS Lambda Database Loader for Amazon Redshift
- Hands-on lab 4: Loading real-time data into an Amazon Redshift database

Module 6: Loading data

- Preparing Data
- Loading data using COPY
- Maintaining tables
- Concurrent write operations
- Troubleshooting load issues
- Hands-on lab 5: Loading data with the COPY command

Day Three

Module 7: Writing queries and tuning for performance

- Amazon Redshift SQL
- User-Defined Functions (UDFs)
- Factors that affect query performance
- The EXPLAIN command and query plans
- Workload Management (WLM)
- Hands-on lab 6: Configuring workload management

Module 8: Amazon Redshift Spectrum

- Amazon Redshift Spectrum
- Configuring data for Amazon Redshift Spectrum
- Amazon Redshift Spectrum Queries
- Hands-on lab 7: Using Amazon Redshift Spectrum

Module 9: Maintaining clusters

- Audit logging
- Performance monitoring
- Events and notifications
- Lab 8: Auditing and monitoring clusters
- Resizing clusters
- Backing up and restoring clusters
- Resource tagging and limits and constraints
- Hands-on lab 9: Backing up, restoring and resizing clusters

Module 10: Analyzing and visualizing data

- Power of visualizations
- Building dashboards
- Amazon QuickSight editions and features

Building Streaming Data Analytics Solutions on AWS

Course description

In this course, you will learn to build streaming data analytics solutions using AWS services, including Amazon Kinesis and Amazon Managed Streaming for Apache Kafka (Amazon MSK). Amazon Kinesis is a massively scalable and durable real-time data streaming service. Amazon MSK offers a secure, fully managed, and highly available Apache Kafka service. You will learn how Amazon Kinesis and Amazon MSK integrate with AWS services such as AWS Glue and AWS Lambda. The course addresses the streaming data ingestion, stream storage, and stream processing components of the data analytics pipeline. You will also learn to apply security, performance, and cost management best practices to the operation of Kinesis and Amazon MSK.

Level	Delivery method	Duration
Intermediate	Instructor-led training, presentations, demonstrations and assessments	1 day

Course objectives

In this course, you will learn to:

- Understand the features and benefits of a modern data architecture. Learn how AWS streaming services fit into a modern data architecture.
- Design and implement a streaming data analytics solution.
- Identify and apply appropriate techniques, such as compression, sharing and partitioning to optimize data storage.
- Select and deploy appropriate options to ingest, transform and store real time and near real time data.
- Choose the appropriate streams, clusters, topics, scaling approach and network topology for a particular business use case.

- Understand how data storage and processing affect the analysis and visualization mechanisms needed to gain actionable business insights.
- Secure streaming data at rest and in transit.
- Monitor analytics workloads to identify and remediate problems.
- Apply cost management best practices.

Intended audience

This course is intended for:
Data Engineers and Architects, Developers who want to build and manage real time applications and streaming data analytics solutions.

Prerequisites

- At least one year of data analytics experience or direct experience building real time applications or streaming analytics solutions. We suggest the Streaming Data Solutions on AWS Whitepaper.

We recommend that attendees of this course have completed:

- Architecting on AWS or Data Analytics Fundamentals and Building Data Lakes on AWS.

Building Streaming Data Analytics Solutions on AWS Course Outline

Module A: Overview of Data Analytics and the Data Pipeline

- Data analytics use cases
- Using the data pipeline for analytics

Module 1: Using Streaming Services in the Data Analytics Pipeline

- The importance of streaming data analytics
- The streaming data analytics pipeline
- Streaming concepts

Module 2: Introduction to AWS Streaming Services

- Streaming data services in AWS
- Amazon Kinesis in analytics solutions
- Demonstration: Explore Amazon Kinesis Data Streams
- Practice Lab: Setting up a streaming delivery pipeline with Amazon Kinesis
- Using Amazon Kinesis Data Analytics
- Introduction to Amazon MSK
- Overview of Spark Streaming

Module 3: Using Amazon Kinesis for Real Time Data Analytics

- Exploring Amazon Kinesis using a clickstream workload
- Creating Kinesis data and delivery streams
- Demonstration: Understanding producers and consumers
- Building Stream Producers
- Building Stream Consumers
- Building and deploying Flink application in Kinesis Data Analytics
- Demonstrations: Explore Zeppelin notebooks for Kinesis Data Analytics
- Practice Lab: Streaming analytics with Amazon Kinesis Data Analytics and Apache Flink

Module 4: Securing, Monitoring and Optimizing Amazon Kinesis

- Optimize Amazon Kinesis to gain actionable business insights
- Security and monitoring best practices

Module 5: Using Amazon MSK in Streaming Data Analytics Solutions

- Use Cases for Amazon MSK
- Creating MSK Clusters
- Demonstration: Provisioning an MSK Cluster
- Ingesting data into Amazon MSK
- Practice Lab: Introduction to access control with Amazon MSK
- Transforming and Processing in Amazon MSK

Module 6: Securing, Monitoring and Optimizing Amazon MSK

- Optimizing Amazon MSK
- Demonstration: Scaling up Amazon MSK storage
- Practice Lab: Amazon MSK streaming pipeline and application deployment
- Security and Monitoring
- Demonstration: Monitoring an MSK Cluster

Module 7: Designing Streaming Data Analytics Solutions

- Use Case Review
- Class Exercise: Designing a streaming data analytics workflow

Module B: Developing Modern Data Architectures on AWS

- Modern Data Architectures

Authoring Visual Analytics Using Amazon QuickSight

Course description

In this course, you will build a data visualization solution using Amazon QuickSight. QuickSight allows everyone in your organization to understand your data by exploring through interactive dashboards, asking questions in natural language, or automatically looking for patterns and outliers powered by machine learning. This course focuses on connecting to data sources, building visuals, designing interactively and creating calculations. You will learn how to apply security best practices to your analyses. You will also explore the machine learning capabilities built into QuickSight.

Level	Delivery method	Duration
Advanced	Instructor-led training, presentations, demonstrations and assessments	2 days

Course objectives

In this course, you will learn to:

- Explain the benefits, use cases and key features of Amazon QuickSight.
- Design, create and customize QuickSight dashboards to visualize data and extract business insights from it.
- Select and configure appropriate visualization types to identify, explore and drill down on the business insights.
- Describe how to use one-click embed to incorporate analytics into applications.
- Connect, transform and prepare data for dashboarding consumption.
- Perform advanced data calculation on QuickSight Analyses.
- Describe the security mechanisms available from Amazon QuickSight.
- Apply fine-grained access control to the dataset.
- Implement machine learning on data sets for anomaly detection and forecasting.
- Explain the benefits and key features of QuickSight Q to enhance the dashboard user experience.

Intended audience

This course is intended for:
Data and Business analysts who build and manage business analytics dashboards

Prerequisites

Students with a minimum of 1 years' experience authoring visual analytics will benefit from this course.

We recommend that attendees of this course have completed:

- Data Analytics Fundamentals

Authoring Visual Analytics Using Amazon QuickSight Course Outline

Module 1: Introduction and Overview of Amazon QuickSight

- Introducing Amazon QuickSight
- Why use Amazon QuickSight for data Visualization

Module 2: Getting Started with Amazon QuickSight

- Interacting with Amazon QuickSight Operational Excellence
- Loading data into Amazon QuickSight
- Visualizing data in Amazon QuickSight
- Demonstration: Walkthrough of Amazon QuickSight interface
- Practice Lab: Create your first dashboard

Module 3: Enhancing and Adding Interactivity to Your Dashboard

- Enhancing your dashboard
- Demonstration: Optimize the size, layout and aesthetics of a dashboard
- Enhancing visualizations with interactivity
- Demonstration: Walkthrough of dashboard interactivity features

- Practice Lab: Enhancing your dashboard

Module 4: Preparing Datasets for Analysis

- Working with datasets
- Demonstration: Transform your datasets for analysis
- Practice Lab: Preparing data for analysis

Module 5: Performing Advanced Data Calculations

- Transform data using advanced calculations
- Practice Lab: Performing advanced data calculations

Activity: Designing a Visual Analytics Solutions

Day Two

Module 6: Overview of Amazon QuickSight Security and Access control

- Overview of Amazon QuickSight security and access control
- Dataset access control in Amazon QuickSight
- Lab: Implementing access control in Amazon QuickSight visualization

Module 7: Exploring machine learning capabilities

- Introducing Machine Learning (ML) insights
- Natural Language Query with QuickSight Q
- Demonstration: Using QuickSight Q
- Lab: Using machine Learning for anomaly detection and forecasting

End of day challenge labs

- Join data sources together
- Create a dashboard
- Enhance the dashboard and add interactivity
- Perform advanced data calculations
- Integrate machine learning tools into the dashboard

Planning and Designing Databases on AWS

Course description

You will learn about the process of planning and designing both relational and non-relational databases. You will learn the design considerations for hosting databases on Amazon Elastic Compute Cloud (Amazon EC2). You will learn about our relational database services including Amazon Relational Database Service (Amazon RDS), Amazon Aurora, and Amazon Redshift. You will also learn about our non-relational database services including Amazon DocumentDB, Amazon DynamoDB, Amazon ElastiCache, Amazon Neptune, and Amazon QLDB. By the end of this course, you will be familiar with the planning and design requirements of all 8 of these AWS databases services, their pros and cons, and how to know which AWS databases service is right for your workloads

Level	Delivery method	Duration
Intermediate	Instructor-led training, hands-on labs, demonstrations, and group exercises	3 days

Course objectives

This course is designed to teach you how to:

- Apply database concepts, database management, and data modeling techniques
- Evaluate hosting databases on Amazon EC2 instances
- Evaluate relational AWS database services and their features (Amazon RDS, Amazon Aurora, and Amazon Redshift)
- Evaluate non-relational AWS database services and their features (Amazon DocumentDB, Amazon DynamoDB, Amazon ElastiCache, Amazon Neptune, and Amazon QLDB)
- Examine how the design criteria apply to each service
- Apply management principles based on the unique features of each service

Intended audience

This course is intended for:

- Data Engineers who are new to designing cloud databases or non-relational databases
- Solutions Architects who are designing services or architectures that are integrated with databases
- Developers that are building cloud database-enabled applications

Prerequisites

We recommend that attendees of this course have:

- Familiarity with AWS Database Services, equivalent to *AWS Database Offerings* digital training
- Understanding of database design concepts, and/or data modeling for relational or non-relational databases
- Familiarity with cloud computing concepts
- Familiarity with general networking and encryption concepts
- Understanding of the three V's of data (volume, velocity, and variety)
- Familiarity with basic data analytics concepts, equivalent to *Data Analytics Fundamentals* digital training
- Understanding of general architecting best practices and the AWS Well-Architected Framework, equivalent to *Architecting on AWS* course

Planning and Designing Databases course outline

Day One

Module 1: Database concepts and general guidelines

- Databases in the cloud
- Database design principles
- Transactional compliance

Module 2: Database planning and design

- Workload requirements
- Design considerations

Module 3: Databases on Amazon EC2

- Amazon EC2 for hosting databases

Module 4: Purpose-built databases

- The journey to AWS
- Data modeling basics

Module 5: Databases on Amazon RDS

- Amazon RDS databases
- Amazon RDS distinguishing features
- Amazon RDS design considerations
- Hands-on Lab 1: working with Amazon RDS databases

Module 6: Databases on Amazon Aurora

- Amazon Aurora databases
- Amazon Aurora distinguishing features
- Amazon Aurora design considerations

Day Two

Module 6: Amazon Aurora *(continued)*

- Hands-on Lab 2: Working with Amazon Aurora databases

Module 7: Amazon DocumentDB (with MongoDB compatibility)

- Amazon DocumentDB
- Amazon DocumentDB design considerations
- Hands-on Lab 3: working with Amazon DocumentDB databases

Module 8: Amazon DynamoDB Tables

- Amazon DynamoDB
- Amazon DynamoDB data modeling
- Amazon DynamoDB distinguishing features
- Amazon DynamoDB design considerations
- Hands-on Lab 4: working with Amazon DynamoDB Tables

Day Three

Module 9: Databases in Amazon Neptune

- Amazon Neptune overview
- Amazon Neptune design considerations

Module 10: Databases in Amazon Quantum Ledger Database (Amazon QLDB)

- Amazon Quantum Ledger Database (Amazon QLDB)
- Amazon QLDB Design Considerations

Module 11: Databases in Amazon ElastiCache

- Amazon ElastiCache
- Amazon ElastiCache for Memcached
- Amazon ElastiCache for Redis

Module 12: Data warehousing in Amazon Redshift

- Amazon Redshift
- Amazon Redshift distinguishing features
- Amazon Redshift data modeling
- Amazon Redshift design considerations
- Hands-on Lab 5: working with Amazon Redshift Clusters

Module 13: Course Overview

Build Applications with AWS NoSQL Databases

Course description

This course is for developers, architects, and database engineers who want to build applications that involve complex data characteristics and millisecond performance requirements from their databases. In this course, you use AWS purpose-built databases to build a typical modern application with diverse access patterns and real-time scaling needs. By the end of the class, you should be able to describe advanced features of Amazon DynamoDB, Amazon DocumentDB (with Mongo compatibility), and Amazon ElastiCache for Redis.

Level	Delivery method	Duration
Intermediate	Presentations, guided tours, group discussions, and hands-on labs.	1 day

Course objectives

In this course you will learn to:

- Build modern applications for the cloud using AWS purpose-built NoSQL databases
- Illustrate solutions using AWS purpose-built databases for handling key-value, document, and in-memory data categories
- Analyze business use cases and apply advanced features of Amazon DynamoDB to implement a scalable solution
- Analyze business use cases and apply advanced features of Amazon ElastiCache to implement a scalable solution
- Analyze business use cases and apply advanced features of Amazon DocumentDB to implement a scalable solution
- Implement event-driven architectures using change streams and AWS Lambda
- Learn how to build solutions faster with Amazon CodeWhisperer

Intended audience

This course is intended for:

- Database Developers
- Solutions Architects
- Database Engineers

Prerequisites

We recommend that attendees of this course have:

- Familiarity with cloud computing concepts
- Familiarity with data modeling for relational or NoSQL databases
- Working experience with Amazon DynamoDB table design
- Working experience with Amazon DocumentDB table design
- Working experience with ElastiCache for Redis
- Familiarity with AWS Lambda and Amazon API Gateway database services
- Familiarity with Python scripting

Build Modern Applications with AWS NoSQL course outline

Module 1: Analyze Use Cases for NoSQL Databases

- Business overview
- Workload solution overview
- AWS NoSQL database portfolio
- Design decisions for a modern application

Module 2: Advanced Amazon DynamoDB Concepts

- Review business workloads for Amazon DynamoDB
- Analyze access patterns and key design
- Create the data model
- Design for performance
- Design event-driven architectures using DynamoDB Streams
- Guided Tour: Design tables using NoSQL Workbench for DynamoDB
- Guided Tour: Use DynamoDB Streams with AWS Lambda
- Hands-on Lab: Implement Fleet and Trip Data Management using Amazon DynamoDB Tables, Indexes, and Change Streams

Module 3: Advanced Amazon DocumentDB Concepts

- Review business workloads for Amazon DocumentDB
- Analyze access patterns
- Create the data model
- Design for performance
- Use Amazon DocumentDB aggregation framework
- Design event-driven architecture using Amazon DocumentDB
- Guided Tour: Document and collection design
- Guided Tour: Aggregation framework
- Guided Tour: Use Amazon DocumentDB Change streams with AWS Lambda
- Hands-on Lab: Implement and Optimize User Profile Data Management Workload on Amazon DocumentDB

Module 4: Advanced Amazon ElastiCache for Redis Concepts

- Review business workloads for Amazon ElastiCache for Redis
- Analyze access patterns

- Create the data model
- Use optimal data structures for the workload
- Guided Tour: Use Amazon ElastiCache for Redis to apply geospatial queries
- Guided Tour: Use Amazon ElastiCache for Redis to natively store and access JSON data
- Guided Tour: Use Amazon ElastiCache for Redis with leaderboards
- Hands-on Lab: Implement Geospatial Bike Searches, User Profile Caching, and Leaderboards with Amazon ElastiCache for Redis

Module 5: Course Summary

- Course review
- AWS Certification levels
- Continue your learning

Developing on AWS

Course description

In this course, you learn how to use the AWS SDK to develop secure and scalable cloud applications using multiple AWS services such as Amazon DynamoDB, Amazon Simple Storage Service, and AWS Lambda. You explore how to interact with AWS using code and learn about key concepts, best practices, and troubleshooting tips.

Level	Delivery method	Duration
Intermediate	Instructor-led training, hands-on labs, and demonstrations	3 days

Course objectives

This course is designed to teach you how to:

- Set up the AWS SDK and developer credentials for Java, C#/.NET and Python
- Interact with AWS services and develop solutions by using the AWS SDK
- Use AWS Identity and Access Management (IAM) for service authentication
- Use Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB as data stores
- Integrate applications and data by using AWS Lambda, Amazon API Gateway, Amazon Simple Queue Service (Amazon SQS), Amazon Simple Notification Service (Amazon SNS), and AWS Step Functions
- Use Amazon Cognito for user authentication
- Use Amazon ElastiCache to improve application scalability
- Leverage the CI/CD pipeline to deploy applications on AWS

Intended audience

This course is intended for:

- Intermediate software developers

Prerequisites

We recommend that attendees of this course have:

- In-depth knowledge of at least one high-level programming language
- Working knowledge of core AWS services and public cloud implementation

Developing on AWS course outline

Day One

Module 1: Introduction to AWS

- Introduction to AWS
- Cloud scenarios
- Infrastructure overview
- Introduction to AWS foundation services

Module 2: Introduction to Developing on AWS

- Getting started developing on AWS
- Introduction to developer tools
- Introduction to management tools

Module 3: Introduction to AWS Identity and Access Management

- Shared responsibility model
- Introduction to IAM
- User authentication and authorization

Module 4: Introduction to the Lab Environment

- Introduction to the lab environment
- Choose your dev instance
- Connect to your dev instance
- Hands-on lab 1: Getting started and working with IAM

Module 5: Developing storage solutions with Amazon Simple Storage Service (Amazon S3)

- Overview of AWS storage options
- Amazon S3 key concepts
- Best practices
- Troubleshooting
- Scenario: Building a complete application
- Hands-on lab 2: Developing storage solutions with Amazon S3

Day Two

Module 6: Developing flexible NoSQL solutions with Amazon DynamoDB

- Introduction to AWS Database Options
- Introduction to Amazon DynamoDB
- Developing with Amazon DynamoDB
- Best practices
- Troubleshooting
- Scenario: Building an end-to end app
- Hands-on lab 3: Developing flexible NoSQL solutions with Amazon DynamoDB

Module 7: Developing Event-Driven solutions with AWS Lambda

- What is serverless computing?
- Introduction to AWS Lambda
  Key concepts
- Use cases
- Best practice
- Scenario: Build and end-to-end app

Module 8: Developing solutions with Amazon API Gateway

- Introduction to Amazon API Gateway
- Developing with Amazon API Gateway
- Best practices
- Introduction to AWS Serverless Application Model (SAM)
- Scenario: Build an end-to-end app
- Hands-on lab 4: Developing Event-Driven Solutions with AWS Lambda

Module 9: Developing solutions with AWS Step Functions

- Understanding the need for AWS Step Functions
- Introduction to AWS Step Functions
- Use cases

Day Three

Module 10: Developing solutions with Amazon Simple Queue Service and Amazon Simple Notification Service

- Why use a queuing service?
- Developing with Amazon Simple Queue Service
- Developing with Amazon Simple Notification Service
- Developing with Amazon MQ
- Hands-on Lab 5: Developing messaging solutions with Amazon SQS and Amazon SNS

Module 11: Caching information with Amazon ElastiCache

- Caching overview
- Caching with Amazon ElastiCache
- Caching strategies

Module 12: Developing Secure Applications

- Securing your applications
- Authenticating your applications to AWS
- Authenticating your customers
- Scenario: Build and end-to-end app

Module 13: Deploying Applications

- Introduction to DevOps
- Introduction to deployment and testing strategies
- Deploying applications with AWS Elastic Beanstalk
- Scenario: Build and end-to-end app
- Hands-on lab 6: Building an end-to-end app

Module 14: Course Wrap-up

- Course Overview
- AWS Training Courses
- Certifications
- Course Feedback

Advanced Developing on AWS

Course description

The Advanced Developing on AWS course uses the real-world scenario of taking a legacy, on-premises monolithic application and refactoring it into a serverless microservices architecture. This three-day advanced course covers advanced development topics such as architecting for a cloud-native environment; deconstructing on-premises, legacy applications and repackaging them into cloud-based, cloud native architectures; and applying the tenets of the Twelve-Factor Application methodology

Level	Delivery method	Duration
Advanced	Instructor-led training, hands-on labs, and group exercises	3 days

Course objectives

In this course, you will learn how to:

- Analyze a monolithic application architecture to determine logical or programmatic break points where the application can be broken up across different AWS services
- Apply Twelve-Factor Application manifesto concepts and steps while migrating from a monolithic architecture
- Recommend the appropriate AWS services to develop a microservices based cloud-native application
- Use the AWS API, CLI, and SDKs to monitor and manage AWS services
- Migrate a monolithic application to a microservices application using the 6 Rs of migration
- Explain the SysOps and DevOps interdependencies necessary to deploy a microservices application in AWS

Intended audience

This course is intended for:

- Experienced software developers who are already familiar with AWS services

Prerequisites

We recommend that attendees of this course have:

- In-depth knowledge of at least one high-level programming language
- Working knowledge of core AWS services and public cloud implementation
- Completion of the *Developing on AWS* course, and then a minimum of 6 months of application of those concepts in a real-world environment

Advanced Developing on AWS course outline

Day One

Module 1: The cloud journey

- Common off-cloud architecture
- Introduction to Cloud Air
- Monolithic architecture
- Migration to the cloud
- Guardrails
- The six R's of migration
- The Twelve-Factor Application Methodology
- Architectural styles and patterns
- Overview of AWS Services
- Interfacing with AWS Services
- Authentication
- Infrastructure as code and Elastic Beanstalk
- Demonstration: Walk through creating base infrastructure with AWS CloudFormation in the AWS console
- Hands-on lab 1: Deploy your monolith application using AWS Elastic Beanstalk

Module 2: Gaining Agility

- DevOps
- CI/CD
- Application configuration
- Secrets management
- CI/CD Services in AWS
- Demonstration: Demo AWS Secrets Manager

Day Two

Module 3: Monolith to Microservices

- Microservices
- Serverless
- A look at Cloud Air
- Microservices using Lambda and API Gateway
- SAM
- Strangling the Monolith
- Hands-on lab: Using AWS Lambda to develop microservices

Module 4: Polyglot Persistence & Distributed Complexity

- Polyglot persistence
- DynamoDB best practices
- Distributed complexity
- Steps functions

Day Three

Module 5: Resilience and Scale

- Decentralized data stores

- Amazon SQS
- Amazon SNS
- Amazon Kinesis Streams
- AWS IoT Message Broker
- Serverless event bus
- Event sourcing and CQRS
- Designing for resilience in the cloud
- Hands-on lab: Exploring the AWS messaging options

Module 6: Security and Observability

- Serverless Compute with AWS Lambda
- Authentication with Amazon Cognito
- Debugging and traceability
- Hands-on lab: Developing microservices on AWS
- Hands-on lab: Automating deployments with CloudFormation

DevOps Engineering on AWS

Course Description

DevOps Engineering on AWS teaches you how to use the combination of DevOps cultural philosophies, practices, and tools to increase your organization's ability to develop, deliver, and maintain applications and services at high velocity on AWS. This course covers Continuous Integration (CI), Continuous Delivery (CD), infrastructure as code, microservices, monitoring and logging, and communication and collaboration. Hands-on labs give you experience building and deploying AWS CloudFormation templates and CI/CD pipelines that build and deploy applications on Amazon Elastic Compute Cloud (Amazon EC2), serverless applications, and container-based applications. Labs for multi-pipeline workflows and pipelines that deploy to multiple environments are also included.

Level	Delivery method	Duration
Intermediate	Instructor-led training, hands-on labs, and group exercises	3 days

Course objectives

In this course, you will learn to:

- Use DevOps best practices to develop, deliver, and maintain applications and services at high velocity on AWS
- List the advantages, roles and responsibilities of small autonomous DevOps teams
- Design and implement an infrastructure on AWS that supports DevOps development projects
- Leverage AWS Cloud9 to write, run and debug your code
- Deploy various environments with AWS CloudFormation
- Host secure, highly scalable, and private Git repositories with AWS CodeCommit
- Integrate Git repositories into CI/CD pipelines
- Automate build, test, and packaging code with AWS CodeBuild
- Securely store and leverage Docker images and integrate them into your CI/CD pipelines

- Build CI/CD pipelines to deploy applications on Amazon EC2, serverless applications, and container-based applications
- Implement common deployment strategies such as “all at once,” “rolling,” and “blue/green”
- Integrate testing and security into CI/CD pipelines
- Monitor applications and environments using AWS tools and technologies

Intended audience

This course is intended for:

- DevOps engineers
- DevOps architects
- Operations engineers
- System administrators
- Developers

Prerequisites

We recommend that attendees of this course have:

- Previous attendance at the Systems Operations on AWS or Developing on AWS courses
- Working knowledge of one or more high-level programing languages, such as C#, Java, PHP, Ruby, Python
- Intermediate knowledge of administering Linux or Windows systems at the command-line level
- Two or more years of experience provisioning, operating, and managing AWS environments

DevOps Engineering on AWS course outline

Day One

Module 1: Introduction to DevOps

- What is DevOps?
- The Amazon journey to DevOps
- Foundations for DevOps

Module 2: Infrastructure automation

- Introduction to Infrastructure Automation
- Diving into the AWS CloudFormation template
- Modifying an AWS CloudFormation template
- Demonstration: AWS CloudFormation template structure, parameters, stacks, updates, importing resources, and drift detection

Module 3: AWS toolkits

- Configuring the AWS CLI
- AWS Software Development Kits (AWS SDKs)
- AWS SAM CLI
- AWS Cloud Development Kit (AWS CDK)
- AWS Cloud9
- Demonstration: AWS CLI and AWS CDK
- Hands-on lab 1: Using AWS CloudFormation to provision and manage a basic infrastructure

Module 4: Continuous integration and continuous delivery (CI/CD) with development tools

- CI/CD Pipeline and Dev Tools

- Demonstration: CI/CD pipeline displaying some actions from AWS CodeCommit, AWS CodeBuild, AWS CodeDeploy and AWS CodePipeline
- Hands-on lab 2: Deploying an application to an EC2 fleet using AWS CodeDeploy

Day Two

Module 5: Introduction to Microservices

- Introduction to Microservices

Module 6: DevOps and containers

- Deploying applications with Docker
- Amazon Elastic Container Service and AWS Fargate
- Amazon Elastic Container Registry and Amazon Elastic Kubernetes service
- Demonstration: CI/CD pipeline deployment in a containerized application

Module 7: DevOps and serverless computing

- AWS Lambda and AWS Fargate
- AWS Serverless Application Repository and AWS SAM
- AWS Step Functions
- Demonstration: AWS Lambda and characteristics
- Demonstration: AWS SAM quick start in AWS Cloud9
- Hands-on lab 3: Deploying a serverless application using AWS Serverless Application Model (AWS SAM) and a CI/CD Pipeline

Module 8: Deployment strategies

- Continuous Deployment
- Deployments with AWS Services

Module 9: Automated testing

- Introduction to testing
- Tests: Unit, integration, fault tolerance, load, and synthetic
- Product and service integrations

Day Three

Module 10: Security automation

- Introduction to DevSecOps
- Security of the Pipeline
- Security in the Pipeline
- Threat Detection Tools
- Demonstration: AWS Security Hub, Amazon GuardDuty, AWS Config, and Amazon Inspector

Module 11: Configuration management

- Introduction to the configuration management process
- AWS services and tooling for configuration management
- Hands-on lab 4: Performing blue/green deployments with CI/CD pipelines and Amazon Elastic Container Service (Amazon ECS)

Module 12: Observability

- Introduction to observability

- AWS tools to assist with observability
- Hands-on lab 5: Using AWS DevOps tools for CI/CD pipeline automations

Module 13: Reference architecture (Optional module)

- Reference architectures

Module 14: Course summary

- Components of DevOps practice
- CI/CD pipeline review
- AWS Certification

MLOps Engineering on AWS

Course description

The course stresses the importance of data, model and code to successful ML deployments. It will demonstrate the use of tools, automation, processes and teamwork in addressing the challenges associated with handoffs between data engineers, data scientists, software developers and operations. The course will also discuss the tools and processes to monitor and act when the model prediction in production starts to drift from agreed-upon key performance indicators.

Level	Delivery method	Duration
Intermediate	Instructor-led training, hands-on labs, and group exercises.	3 days

Course objectives

In this course, you will learn to:

- Describe machine learning operations
- Understand the key differences between DevOps and MLOps
- Describe the machine learning workflow
- Discuss the importance of communications in MLOps
- Explain end-to-end options for automation of ML workflows
- List key Amazon SageMaker features for MLOps automation
- Build an automated ML Process that builds, trains, tests and deploys models
- Build an automated ML Process that retrains the model based on change(s) to the model code
- Identify elements and important steps in the deployment process
- Describe items that might be included in a model package and their use in training and inference
- Recognize Amazon SageMaker options for detecting models for deployment, including support for ML Frameworks and built-in algorithms or bring-your-models
- Differentiate scaling in Machine Learning from scaling in other applications
- Determine when to use different approaches to inference
- Discuss deployment strategies, benefits, challenges and typical use cases
- Describe the challenges when deploying machine learning to edge devices
- Recognize important Amazon SageMaker features that are relevant to deployment and inference
- Describe why monitoring is important
- Detect data drifts in the underlying input data
- Demonstrate how to monitor model resource consumption and latency

- Discuss how to integrate human-in-the-loop reviews of model results in production

Intended audience

This course is intended for:

- DevOps engineers
- ML Engineers
- Developers/operations with responsibility for operationalizing ML Models

Prerequisites

Required:

- Completed the AWS Technical Essentials course
- *DevOps Engineering on AWS* course or equivalent experience
- Practical Data Science with Amazon SageMaker course or equivalent

Recommended:

- *The Elements of Data Science* (Digital Course) or equivalent experience
- Machine Learning Terminology and Process (Digital Course)

MLOps Engineering on AWS Course outline

Day One

Module 1: Introduction to MLOps

- Machine Learning Operations
- Goals of MLOps
- Communication
- From DevOps to MLOps
- ML Workflow
- Scope
- MLOps view of ML Workflow
- MLOps cases

Module 2: MLOps Development

- Intro to build, train and evaluate machine learning models
- MLOps security
- Automating
- Apache AirFlow
- Kubernetes integration to MLOps
- Amazon SageMaker for MLOps
- Lab 1: Bring your own algorithm to an MLOp pipeline
- Demonstration: Amazon SageMaker
- Intro to build, train and evaluate machine learning models
- Lab 2: Code and Serve your ML model and AWS CodeBuild
- Activity: MLOps Action Plan Workbook

Day Two

Module 3: MLOps Deployment

- Introduction to deployment operations
- Model packaging

- Inference
- Lab 3: Deploy your model to production
- SageMaker production variants
- Deployment Strategies
- Deploying to the edge
- Lab 4: Conduct A/B testing
- Activity: MLOps Action Plan Workbook

Day Three

Module 4: Model Monitoring and Operations

- Lab 5: Troubleshoot your pipeline
- The importance of monitoring
- Monitoring by design
- Lab 6: Monitor your ML model
- Human-in-the-loop
- Amazon SageMaker Model Monitor
- Demonstration: Amazon SageMaker pipelines, Model Monitor, model registry and feature store
- Solving the problem(s)
- Activity: MLOps Action Plan Workbook

Module 5: Wrap-Up

- Course review
- Activity: MLOps Action Plan Workbook
- **Wrap-up**

Practical Data Science with Amazon SageMaker

Course description

You will learn how to solve a real-world use case with Machine Learning (ML) and produce actionable results using Amazon SageMaker. This course walks through the stages of a typical data science process for Machine Learning from analyzing and visualizing a dataset to preparing the data, and feature engineering. Individuals will also learn practical aspects of model building, training, tuning, and deployment with Amazon SageMaker. Real life use case includes customer retention analysis to inform customer loyalty programs.

Level	Delivery method	Duration
Intermediate	Instructor-led training, hands-on labs, and group exercises	1 day

Course objectives

In this course, you will learn how to:

- Prepare a dataset for training
- Train and evaluate a Machine Learning model
- Automatically tune a Machine Learning model
- Prepare a Machine Learning model for production
- Think critically about Machine Learning model results

Intended audience

This course is intended for:

- Developers
- Data Scientists

Prerequisites

- Familiarity with Python programming language
- Basic understanding of Machine Learning

Practical Data Science with Amazon SageMaker course outline

Day One

Module 1: Introduction to machine learning

- Types of ML
- Job Roles in ML
- Steps in the ML pipeline

Module 2: Introduction to data prep and SageMaker

- Training and test dataset defined
- Introduction to SageMaker
- Demonstration: SageMaker console
- Demonstration: Launching a Jupyter notebook

Module 3: Problem formulation and dataset preparation

- Business challenge: Customer churn
- Review customer churn dataset

Module 4: Data analysis and visualization

- Demonstration: Loading and visualizing your dataset
- Exercise 1: Relating features to target variables
- Exercise 2: Relationships between attributes
- Demonstration: Cleaning the data

Module 5: Training and evaluating a model

- Types of algorithms
- XGBoost and SageMaker
- Demonstration: Training the data
- Exercise 3: Finishing the estimator definition
- Exercise 4: Setting hyper parameters
- Exercise 5: Deploying the model
- Demonstration: hyper parameter tuning with SageMaker
- Demonstration: Evaluating model performance

Module 6: Automatically tune a model

- Automatic hyper parameter tuning with SageMaker
- Exercises 6-9: Tuning jobs

Module 7: Deployment / production readiness

- Deploying a model to an endpoint
- A/B deployment for testing
- Auto Scaling

- Demonstration: Configure and test auto scaling
- Demonstration: Check hyper parameter tuning job
- Demonstration: AWS Auto Scaling
- Exercise 10-11: Set up AWS Auto Scaling

Module 8: Relative cost of errors

- Cost of various error types
- Demo: Binary classification cutoff

Module 9: Amazon SageMaker architecture and features

- Accessing Amazon SageMaker notebooks in a VPC
- Amazon SageMaker batch transforms
- Amazon SageMaker Ground Truth
- Amazon SageMaker Neo

Amazon SageMaker Studio for Data Scientists

Course description

Amazon SageMaker Studio helps data scientists prepare, build, train, deploy, and monitor machine learning (ML) models quickly by bringing together a broad set of capabilities purpose-built for ML. This course prepares experienced data scientists to use the tools that are part of SageMaker Studio to improve productivity at every step of the ML lifecycle.

Level	Delivery method	Duration
Advanced	Instructor-led training, presentations, demonstrations and assessments	3 days

Course objectives

In this course, you will learn to:

- Accelerate the preparation, building, training, deployment and monitoring of ML solutions using Amazon SageMaker Studio

Intended audience

This course is intended for:
Experience Data Scientist who are proficient in ML and deep learning fundamentals. Relevant experience includes using ML Frameworks, Python programming and the process of building, training, tuning and deploying models.

Prerequisites

- At least one year of data analytics experience or direct experience building real time applications or streaming analytics solutions. We suggest the Streaming Data Solutions on AWS Whitepaper.

We recommend that attendees of this course have completed:

- AWS Tech Essentials 1-day instructor led course

Amazon SageMaker Studio for Data Scientists Course Outline

Module 1: Amazon SageMaker Setup and Navigation

- Launch SageMaker Studio from the AWS Service Catalog.
- Navigate the SageMaker Studio UI.
- Demo 1: SageMaker UI Walkthrough
- Lab 1: Launch SageMaker Studio from AWS Service Catalog

Module 2: Data Processing

- Use Amazon SageMaker Studio to collect, clean, visualize, analyze, and transform data.
- Set up a repeatable process for data processing.
- Use SageMaker to validate that collected data is ML ready.
- Detect bias in collected data and estimate baseline model accuracy.
- Lab 2: Analyze and Prepare Data Using SageMaker Data Wrangler
- Lab 3: Analyze and Prepare Data at Scale Using Amazon EMR
- Lab 4: Data Processing Using SageMaker Processing and the SageMaker Python SDK
- Lab 5: Feature Engineering Using SageMaker Feature Store

Module 3: Model Development

- Use Amazon SageMaker Studio to develop, tune, and evaluate an ML model against business objectives and fairness and explainability best practices.
- Fine-tune ML models using automatic hyperparameter optimization capability.
- Use SageMaker Debugger to surface issues during model development.
- Demo 2: Autopilot
- Lab 6: Track Iterations of Training and Tuning Models Using SageMaker Experiments
- Lab 7: Analyze, Detect, and Set Alerts Using SageMaker Debugger
- Lab 8: Identify Bias Using SageMaker Clarify

Module 4: Deployment and Inference

- Use Model Registry to create a model group; register, view, and manage model versions; modify model approval status; and deploy a model.
- Design and implement a deployment solution that meets inference use case requirements.
- Create, automate, and manage end-to-end ML workflows using Amazon SageMaker Pipelines.
- Lab 9: Inferencing with SageMaker Studio
- Lab 10: Using SageMaker Pipelines and the SageMaker Model Registry with SageMaker Studio

Module 5: Monitoring

- Configure a SageMaker Model Monitor solution to detect issues and initiate alerts for changes in data quality, model quality, bias drift, and feature attribution (explainability) drift.
- Create a monitoring schedule with a predefined interval.
- Demo 3: Model Monitoring

Module 6: Managing SageMaker Studio Resources

- List resources that accrue charges.
- Recall when to shut down instances.
- Explain how to shut down instances, notebooks, terminals, and kernels.
- Understand the process to update SageMaker Studio.

Capstone

- The Capstone lab will bring together the various capabilities of SageMaker Studio discussed in this course. Students will be given the opportunity to prepare, build, train, and deploy a model using a tabular dataset not seen in earlier labs. Students can choose among basic, intermediate, and advanced versions of the instructions.
- Capstone Lab: Build an End-to-End Tabular Data ML Project Using SageMaker Studio and the SageMaker Python SDK

Developing Generative AI Applications on AWS

Course description

This course is designed to introduce generative AI to software developers interested in leveraging large language models without fine-tuning. The course provides an overview of generative AI, planning a generative AI project, getting started with Amazon Bedrock, the foundations of prompt engineering, and the architecture patterns to build generative AI applications using Amazon Bedrock and LangChain.

Level	Duration	Format	Delivery method
Advanced	2 days	Presentations, demonstrations and group exercises	Instructor-led training

Course objectives

This course is you will learn to:

- Describe generative AI and how it aligns to machine learning
- Define the importance of generative AI and explain its potential risks and benefits
- Identify business value from generative AI use cases
- Discuss the technical foundations and key terminology for generative AI
- Explain the steps for planning a generative AI project
- Identify some of the risks and mitigations when using generative AI
- Understand how Amazon Bedrock works
- Familiarize yourself with basic concepts of Amazon Bedrock
- Recognize the benefits of Amazon Bedrock
- List typical use cases for Amazon Bedrock
- Describe the typical architecture associated with an Amazon Bedrock solution
- Understand the cost structure of Amazon Bedrock
- Implement a demonstration of Amazon Bedrock in the AWS Management Console
- Define prompt engineering and apply general best practices when interacting with FMs
- Identify the basic types of prompt techniques, including zero-shot and few-shot learning
- Apply advanced prompt techniques when necessary for your use case
- Identify which prompt-techniques are best-suited for specific models
- Identify potential prompt misuses
- Analyze potential bias in FM responses and design prompts that mitigate that bias
- Identify the components of a generative AI application and how to customize a foundation model (FM)

- Describe Amazon Bedrock foundation models, inference parameters, and key Amazon Bedrock APIs
- Identify Amazon Web Services (AWS) offerings that help with monitoring, securing, and governing your Amazon Bedrock applications
- Describe how to integrate LangChain with large language models (LLMs), prompt templates, chains, chat models, text embeddings models, document loaders, retrievers, and Agents for Amazon Bedrock
- Describe architecture patterns that can be implemented with Amazon Bedrock for building generative AI applications
- Apply the concepts to build and test sample use cases that leverage the various Amazon Bedrock models, LangChain, and the Retrieval Augmented Generation (RAG) approach

Intended audience

This course is intended for software developers interested in leveraging large language models without fine-tuning

Prerequisites

We recommend that attendees of this course have completed:

- AWS Technical Essentials
- Intermediate-level proficiency in Python

Developing Generative AI Applications on AWS Course outline

Day One

Module 1: Introduction to Generative AI - Art of the Possible

- Overview of ML
- Basics of generative AI
- Generative AI use cases
- Generative AI in practice
- Risks and benefits

Module 2: Planning a Generative AI Project

- Generative AI fundamentals
- Generative AI in practice
- Generative AI context
- Steps in planning a generative AI project
- Risks and mitigation

Module 3: Getting Started with Amazon Bedrock

- Introduction to Amazon Bedrock
- Architecture and use cases
- How to use Amazon Bedrock
- Demonstration: Setting Up Bedrock Access and Using Playgrounds

Module 4: Foundations of Prompt Engineering

- Basics of foundation models
- Fundamentals of prompt engineering
- Basic prompt techniques
- Advanced prompt techniques

- Demonstration: Fine-Tuning a Basic Text Prompt
- Model-specific prompt techniques
- Addressing prompt misuses
- Mitigating bias
- Demonstration: Image Bias-Mitigation

Day Two

Module 5: Amazon Bedrock Application Components

- Applications and use cases
- Overview of generative AI application components
- Foundation models and the FM interface
- Working with datasets and embeddings
- Demonstration: Word Embeddings
- Additional application components
- RAG
- Model fine-tuning
- Securing generative AI applications
- Generative AI application architecture

Module 6: Amazon Bedrock Foundation Models

- Introduction to Amazon Bedrock foundation models
- Using Amazon Bedrock FMs for inference
- Amazon Bedrock methods
- Data protection and auditability
- Demonstration: Invoke Bedrock Model for Text Generation Using Zero-Shot Prompt

Module 7: LangChain

- Optimizing LLM performance
- Integrating AWS and LangChain
- Using models with LangChain
- Constructing prompts
- Structuring documents with indexes
- Storing and retrieving data with memory
- Using chains to sequence components
- Managing external resources with LangChain agents
- Demonstration: Bedrock with LangChain Using a Prompt that Includes Context

Module 8: Architecture Patterns

- Introduction to architecture patterns
- Text summarization
- Demonstration: Text Summarization of Small Files with Anthropic Claude
- Demonstration: Abstractive Text Summarization with Amazon Titan Using LangChain
- Question answering
- Demonstration: Using Amazon Bedrock for Question Answering
- Chatbots
- Demonstration: Conversational Interface – Chatbot with AI21 LLM
- Code generation
- Demonstration: Using Amazon Bedrock Models for Code Generation

- LangChain and agents for Amazon Bedrock
- Demonstration: Integrating Amazon Bedrock Models with LangChain Agents

AWS Migration Essentials

Course description

This course is intended to provide solution architects with the foundational knowledge required to successfully plan and perform lift and shift migrations to the AWS Cloud. In this course you will learn about methodologies for discovering, performing and tracking migrations using various AWS Tools and Services.

Level	Duration	Format	Delivery method
Fundamental	1 day	Presentations, demonstrations and group exercises	Instructor-led training

Course objectives

This course is you will learn to:

- Determine cloud readiness and migration strategies using assessment tools and services provided by Amazon Web Services (AWS)
- Articulate the key tasks involved in planning and mobilizing migrations
- Describe at high-level, the Amazon Web Services (AWS) services, resources and tools necessary for Migrations of data and databases
- Describe at high-level, the Amazon Web Services (AWS) services, resources and tools necessary for Migrations of applications.

Intended audience

This course is intended for a technical audience (Solutions Architects, Developers and Administrators) with limited or no knowledge on cloud migration.

Prerequisites

We recommend that attendees of this course have completed:

- AWS Technical Essentials

Course outline

Module 0: Introduction

- Introductions
- Course Overview

Module 1: Assess

- Migration Phases
- Migration drivers and outcomes
- Cloud Adoption Readiness Tool (CART)
- Migration Readiness Assessment (MRA)

- Migration Evaluator
- Migration Portfolio Assessment (MPA)

Module 2: Mobilize

- Landing Zones
- AWS Application Discovery Service
- Migration Strategies
- AWS Migration Hub

Module 3: Migrate and Modernize: Database and Data Migration

- AWS Database Migration Service (AWS DMS)
- Data Migration
- Lab 1: Database Migration with AWS DMS

Module 4: Migrate and Modernize: AWS Application Migration Service

- Migrate Servers with AWS Application Migration Service (AWS MGN)
- Modernization Phases
- AWS Well-Architected Framework for Migration
- Application Optimization
- Lab 2: Application Migration with AWS MGN

Module 5: Course Summary

- Course Summary

Migrating to AWS

Course description

This course is for individuals who seek an understanding of how to plan and migrate existing workloads to the AWS Cloud. You learn about various cloud migration strategies and how to apply each step of the migration process, including portfolio discovery, application migration planning and design, performing a migration and post-migration validation and application optimization. Hand-on labs reinforce learning and each lab is designed to provide you with the foundation necessary to complete migration tasks in your organization.

Level	Delivery method	Duration
Intermediate	Instructor-led training, hands-on labs, and group exercises	3 days

Course objectives

This course is designed to teach you how to:

- Recognize the common business and technical drivers for migrating to the cloud
- Summarize the three phases of a migration and associated objectives, tasks, and stakeholders for each
- Describe AWS architecture, tools, and migration best practices
- Distinguish between the various cloud migration strategies and when each is most appropriate
- Determine an organizations application migration readiness

- Discover a portfolio and gather data necessary for migration
- Plan and design and application migration strategy
- Perform and validate application migration to the cloud
- Optimize applications and operations after migrating to the cloud

Intended audience

This course is intended for:

- Solutions Architects and Engineers who perform cloud migrations
- IT Project Managers who are involved in projects related to migrating existing workloads to the AWS Cloud
- Operations leads

Prerequisites

We recommend that attendees of this course have:

- Taken *AWS Technical Essentials* or *Architecting on AWS* course
- Familiarity with enterprise IT infrastructure (hardware and software)

Migrating to AWS course outline

Day One

Module 1: What is a Cloud Migration?

- Drivers and outcomes of a cloud migration
- Planning for a successful cloud migration
- The three-phase migration process

Module 2: Assessing Migration Readiness

- The assess phase
- Cloud readiness assessment tools
- Examine your IT landscape and build your business case
- Group Exercise: The Cloud Adoption Tool (CART)

Module 3: Preparing for a Migration and Understanding Related Workstreams

- Mobilize phase
- Migration-related workstreams

Module 4: Discovering Landing Zones and Their Benefits

- What is a landing zone?
- Custom multi-account structure with AWS Organizations
- AWS Control Tower
- Customizations for AWS Control Tower (CfCt)
- Planning for connectivity

Module 5: Building a Landing Zone

- Planning a landing zone
- Design a multi-account structure
- Governance polices
- Demonstration: AWS Control Tower

- Hands-On Lab: Connecting Your On-Premises Network and Directory Services to AWS

Module 6: Discovering the Portfolio and Understanding Migration Strategies

- Detailed portfolio discovery workstream
- Evaluating cloud readiness
- Cloud migration strategies
- Group Exercise: Choose a migration strategy (scenario-based)

Day Two

Module 7: Understanding and Choosing Portfolio Discovery Tools

- Migration Evaluator
- AWS Migration Hub and AWS Application Discovery Service
- AWS Systems Manager and Amazon CloudWatch
- Hands-On Lab: Gathering Application Data Necessary for Migration

Module 8: Planning and Designing a Migration

- Plan the migration overall
- Building the migration factory
- Design the migration for each application
- Group Exercise: Build a migration plan
- Group Exercise: Design for migration

Module 9: Performing the Migrations to AWS

- Server migration process
- Server migration tools
- VMware Cloud on AWS
- AWS Application Migration Service (AWS MGN)
- Evaluating server migration tools
- Hands-On Lab: Migrating an Application to AWS

Module 10: Understanding Database and Data Migration Services

- Data migration
- Online data migration services
- Offline data migration services
- Database migration
- Hands-On Lab: Migrating an Existing Database to Amazon Aurora

Day Three

Module 11: Understanding Additional Migration Support Options

- AWS Managed Services
- AWS Service Catalog
- AWS Service Catalog integrations
- Microsoft workloads on AWS
- SAP on AWS

Module 12: Integrating, Validating, and Cutting Over Applications

- Migrate and modernize phase
- Cutover strategy

Module 13: Modernizing and Optimizing Applications

- Cost optimization
- Performance optimization
- AWS tools used to optimize
- Modernize the enterprise
- Use containers
- Use serverless architectures
- Hands-On Lab: Optimizing an Application with Amazon S3 and Amazon ECS

Module 14: Understanding Operations Tools, Integration Testing, and Automation

- AWS Config
- Infrastructure and operations as code
- Adopting a DevOps approach
- Automate change and configuration
- Automate management

Module 15: Migration Best Practices

- Course review
- Best practices
- Continue your learning
- Hands-On Lab: Automating Application Deployments

Networking Essentials for Cloud Applications on AWS

Course description

This course provides a comprehensive understanding of networking concepts and services withing Amazon Web Services Cloud environment. Designed for novice and experiences network engineers, this course covers essential topics, best practices, and hands-on labs. Its purpose is to equip learners with the knowledge and skills that are required to design, configure and optimize network infrastructure on AWS.

Level	Delivery method	Duration
Intermediate	Instructor-led training, presentations, demonstrations and assessments	1 day

Course objectives

In this course, you will learn to:

- Design a network infrastructure for a scalable production application, considering design tradeoffs between different networking services.
- Configure networking services for a highly available, resilient and scalable application.
- Implement the networking infrastructure according to evolving business requirements

- Implement Networking best practices to align towards AWS Well-Architected Framework

Intended audience

This course is intended for:

- Newly hired cloud engineers
- On-premise IT Engineers
- Cloud Architects
- Cloud Engineers
- Network Engineers

Prerequisites

Basic knowledge of networking concepts and AWS Services

We recommend that attendees of this course have completed:

- AWS Security Essentials or AWS Cloud Practitioner Essentials

Networking Essentials for Cloud Applications on AWS Course Outline

Module 1: Networking on AWS

- IP Addressing
- Amazon Virtual Private Cloud (Amazon VPC) fundamentals
- Subnets
- Amazon VPC IP Address Manager (IPAM)
- Elastic Network Interfaces
- Elastic IP Addressing
- Route Table
- Internet and NAT gateways
- Basic Traffic Filtering mechanisms for VPC
- Knowledge Check

Module 2: Load Balancing and Scaling on AWS

- Elastic Load Balancing (ELB)
- Cross Zone Load Balancing
- Auto Scaling Group (ASG) basics
- Knowledge Check
- Use Case Part one
- Hands-on lab: Building a multi-Availability Zone VPC Architecture

Module 3: VPC Interconnectivity and Content Delivery

- VPC Interconnectivity
- VPC Peering
- VPC Transit Gateway
- VPC Endpoints
- Edge Locations
- AWS Global Accelerator
- Knowledge Check

- Use Case Part Two
- Hands-on lab: Accelerating Performance with Amazon CloudFront

Module 4: High Availability with Amazon Route 53

- Amazon Route 53
- Knowledge Check
- Use Case Part 3
  Hands-on lab: Achieving Fault Tolerance and Global Traffic Optimization

Module 5: Course Wrap Up

- Course Reflection, labs recap, use case conclusion and feedback.

Cloud Operations on AWS

Course description

This course teaches systems operators and anyone performing system operations functions how to install, configure, automate, monitor, secure, maintain and troubleshoot the services, networks, and systems on AWS necessary to support business applications. The course also covers specific AWS features, tools, and best practices related to these functions

Level	Delivery method	Duration
Intermediate	Instructor-led training, hands-on labs, and group exercises	3 days

Course objectives

This course is designed to teach you how to:

- Identify the AWS services that support the different phases of Operational Excellence, a Well-Architected Framework Pillar
- Manage access to AWS resources use AWS accounts and Organizations and AWS Identity and Access Management (AWS IAM)
- Maintain an inventory of in-use AWS resources using AWS services such as AWS Systems Manager, AWS CloudTrail and AWS Config
- Develop a resource deployment strategy utilizing metadata tags, Amazon Machine Images and Control Tower to deploy and maintain and AWS Cloud environment
- Automate resource deployment using AWS services such as AWS CloudFormation and AWS Service Catalog
- Use AWS Services to manage AWS resources through CloudOps lifecycle processes such as deployments and patches
- Configure a highly available cloud environment that leverages AWS services such as Amazon Route 53 and Elastic Load Balancing to route traffic for optimal latency and performance
- Configure AWS Auto Scaling and Amazon EC2 auto scaling to scale your cloud environment based on demand
- Use Amazon CloudWatch and associated features such as alarms, dashboards, and widgets to monitor your cloud environment
- Manage permissions and track activity in your cloud environment using AWS services such as AWS CloudTrail and AWS Config

- Deploy your resources to an Amazon Virtual Private Cloud (Amazon VPC), establish necessary connectivity to your Amazon VPC, and protect your resources from disruptions of service.
- State the purpose, benefits, and appropriate use cases for mountable storage in your AWS cloud environment

Intended audience

This course is intended for:

- Systems administrators and operators who are operating in the AWS Cloud
- Informational technology workers who want to increase their systems operations knowledge

Prerequisites

We recommend that attendees of this course have:

- Successfully completed the ***AWS Technical Essentials*** course
- A background in either software development or systems administration
- Proficiency in maintaining operating systems at the command line, such as shell scripting in Linux environments or cmd/PowerShell in Windows
- Basic knowledge of networking protocols (TCP/IP, HTTP)

Cloud Operations on AWS course outline

Day One

Module 1: Introduction to Cloud Operations on AWS

- What is Cloud Operations
- AWS Well-Architected Framework
- AWS Well-Architected Tool

Module 2: Access Management

- AWS Identity and Access Management (IAM)
- Resources, accounts and AWS Organizations

Module 3: Systems Discovery

- Methods to interact with AWS services
- Tools for automating resource discovery
- Inventory with AWS Systems Manager and AWS Config
- Hands-on lab: Auditing AWS Resources with AWS Systems Manager and AWS Config

Module 4: Deploy and Update Resources

- Cloud operations in deployments
- Tagging Strategies
- Deployment with Amazon Machine Images (AMI's)
- Deployment using AWS Control Tower

Module 5: Automate Resource Deployment

- Deployment using AWS CloudFormation
- Deployment using AWS Service Catalog
- Hands-on lab: Infrastructure as Code

Day Two

Module 6: Manage Resources

- AWS Systems Manager
- Hands-on lab: Operations as Code

Module 7: Configure Highly Available Systems

- Distributing Traffic with Elastic Load Balancing
- Amazon Route 53

Module 8: Automate Scaling

- Scaling with AWS Auto Scaling
- Scaling with Spot Instances
- Managing licenses with AWS License Manager

Module 9: Monitoring and Maintaining System Health

- Monitoring and maintaining healthy workloads
- Monitoring AWS Infrastructure
- Monitoring applications
- Hands-on lab: Monitoring Applications and Infrastructure

Module 10: Data Security and System Auditing

- Maintaining a strong identity and access foundation
- Implementing detection mechanisms
- Automating incident remediation

Day Three

Module 11: Operate Secure Resilient Networks

- Building a secure Amazon Virtual Private Network (Amazon VPC)
- Networking beyond the VPC

Module 12: Mountable Storage

- Configuring Amazon Elastic Block Storage (Amazon EBS)
- Sizing Amazon EBS volumes for performance
- Using Amazon EBS snapshots
- Using Amazon Data Lifecycle Manager to manager your AWS resources
- Creating backup and data recovery plans
- Configuring shared file system storage
- Hands-on lab: Automating with AWS Backup for Archiving and Recovery

Module 13: Object Storage

- Deploying Amazon Simple Storage Service (Amazon S3)
- Managing storage lifecycles on Amazon S3

Module 14: Cost Reporting, Alerts and Optimization

- Gaining AWS cost awareness
- Using control mechanisms for cost management

- Optimizing your AWS Spend and Usage
- Hands-on lab: Capstone lab for CloudOps

Developing Serverless Solutions on AWS

Course description

This course gives developers exposure to and practice with best practices for building serverless applications using AWS Lambda and other services in the AWS serverless platform. You will use AWS frameworks to deploy a serverless application in hands-on labs that progress from simpler to more complex topics. You will use AWS documentation throughout the course to develop authentic methods for learning and problem-solving beyond the classroom

Level	Delivery method	Duration
Intermediate	Instructor-led training, hands-on labs, and demonstrations	3 days

Course objectives

In this course, you will learn to:

- Apply event-driven best practices to a serverless application design using appropriate AWS services
- Identify the challenges and trade-offs of transitioning to serverless development, and make recommendations that suit your development organization and environment
- Build serverless applications using patterns that connect AWS managed services together, and account for service characteristics, including service quotas, available integrations, invocation model, error handling, and event source payload
- Compare available options for writing infrastructure as code, including AWS CloudFormation, AWS Amplify, AWS Serverless Application Model (AWS SAM), and AWS Cloud Development Kit (AWS CDK)
- Apply best practices to writing Lambda functions inclusive of error handling, logging, environment re-use, using layers, statelessness, idempotency, and configuring concurrency and memory
- Apply best practices for building observability and monitoring into your serverless application
- Apply security best practices to serverless applications
- Identify key scaling considerations in a serverless application, and match each consideration to the methods, tools, or best practices to manage it
- Use AWS SAM, AWS CDK, and AWS developer tools to configure a CI/CD workflow, and automate deployment of a serverless application
- Create and actively maintain a list of serverless resources that will assist in your ongoing serverless development and engagement with the serverless community

Intended audience

This course is intended for:

- Developers who have some familiarity with serverless and experience with development in the AWS Cloud

Prerequisites

We recommend that attendees of this course have:

- Familiarity with the basics of AWS Cloud Architecture
- An understanding of developing applications on AWS equivalent to completing the *Developing on AWS* course
- Knowledge equivalent to completing the following serverless digital trainings: *AWS Lambda Foundations* and *Amazon API Gateway for Serverless Applications*

Developing Serverless Solutions on AWS Course Outline

Day One

Module 1: Thinking Serverless

- Best practices for building modern serverless applications
- Event-driven design
- AWS services that support event-driven serverless applications

Module 2: API-Driven Development and Synchronous Event Sources

- Characteristics of standard request/response API-based web applications
- How Amazon API Gateway fits into serverless applications
- Try it out exercise: Set up an HTTP API endpoint integrated with a Lambda function
- High-level comparison of API types (REST/HTTP, WebSocket, GraphQL)

Module 3: Introduction to Authentication, Authorization and Access Control

- Authentication vs. Authorization
- Options for authentication to APIs using API Gateway
- Amazon Cognito in serverless applications
- Amazon Cognito user pools vs. federated identities

Module 4: Serverless Deployment Frameworks

- Overview of imperative vs. declarative programming for infrastructure as code
- Comparison of CloudFormation, AWS CDK, Amplify and AWS SAM frameworks
- Features of Aws SAM and the AWS SAM CLI for local emulation and testing

Module 5: Using Amazon EventBridge and Amazon SNS to Decouple Components

- Development considerations when using asynchronous event sources
- Features and use cases of Amazon EventBridge
- Try it out exercise: Build a custom EventBridge bus and rule
- Comparison of use cases for Amazon Simple Notification Service (Amazon SNS) vs. EventBridge
- Try it out exercise: Configure and Amazon SNS topic with filtering

Module 6: Event-Driven Development using Queues and Streams

- Development considerations when using polling event sources to trigger Lambda functions
- Distinctions between queues and streams as event sources for Lambda
- Selecting appropriate configurations when using Amazon Simple Queue Service (Amazon SQS) or Amazon Kinesis Data Streams as an event source for Lambda

- Try it out exercise: Configure an Amazon SQS queue with a dead-letter queue as a Lambda event source

Hands-On Labs

- Hands-on Lab 1: Deploying a Simple Serverless Application
- Hands-on Lab 2: Message Fan-Out with Amazon EventBridge

Day Two

Module 7: Writing Good Lambda Functions

- How the Lambda lifecycle influences your function mode
- Best practices for your lambda functions
- Configuring a function
- Function code, versions and aliases
- Try it out exercise: Configure and test a Lambda function
- Lambda error handling
- Handling partial failures with queues and streams

Module 8: Step Functions for Orchestration

- AWS Step Functions in serverless architectures
- Try it out exercise: Step functions states
- The call back pattern
- Standard vs. Express workflows
- Step functions direct integrations
- Try it out exercise: Troubleshooting a standard step function workflow

Module 9: Observability and Monitoring

- The three pillars of observability
- Amazon CloudWatch Logs insights
- Writing effective log files
- Try it out exercise: interpreting logs
- Using AWS X-Ray for observability
- Try it out exercise: Enable X-Ray and interpret X-Ray traces
- CloudWatch Metrics and embedded metrics format
- Try it out exercise: Metrics and alarms
- Try it out exercise: Service Lens

Hands-On Labs

- Hands-on Lab 3: Workflow Orchestration using AWS Step Functions
- Hands-on Lab 4: Observability and Monitoring

Day Three

Module 10: Serverless Application Security

- Security best practices for serverless applications
- Applying security at all layers
- API Gateway and application security
- Lambda and application security
- Protecting data in your serverless data stores
- Auditing and traceability

Module 11: Handling Scale in Serverless Applications

- Scaling consideration for serverless applications
- Using API Gateway to manage scale
- Lambda concurrency scaling
- How different event sources scale with Lambda

Module 12: Automating the Deployment Pipeline

- The importance of CI/CD in serverless applications
- Tools in a serverless pipeline
- AWS SAM features for serverless deployment
- Best practices for automation
- Course Wrap-up

Hands-on Labs:

- Hands-on Lab 5: Securing Serverless applications
- Hands-on Lab 6: Serverless CS/CD on AWS

AWS Security Essentials

Course description

This course covers fundamental AWS cloud security concepts, including AWS access control, data encryption methods, and how network access to your AWS infrastructure can be secured. Based on the AWS Shared Security Model, you learn where you are responsible for implementing security in the AWS Cloud and what security-oriented services are available to you and why and how the security services can help meet the security needs of your organization.

Level	Duration	Format	Delivery method
Fundamental	1 day	Instructor-led training and hands-on labs	Classroom, virtual classroom, and private on-site training.

Course objectives

This course is designed to teach you how to:

- Identify security benefits and responsibilities when using the AWS Cloud
- Describe the access control and management features of AWS
- Understand the different data encryption methods to secure sensitive data
- Describe how to secure network access to your AWS resources
- Determine which AWS services can be used for security logging and monitoring

Intended audience

This course is intended for:

- IT business-level professionals interested in cloud security practices
- Security professionals with minimal working knowledge of AWS

Prerequisites

We recommend that attendees of this course have:

- Working knowledge of IT Security practices and infrastructure concepts, familiarity with cloud computing concepts

AWS Security Essentials Course outline

Module 1: Security on AWS

- Security design principles in the AWS Cloud
- AWS Shared Responsibility Model

Module 2: Security OF the Cloud

- AWS Global Infrastructure
- Data Center Security
- Compliance and Governance

Module 3: Security IN the Cloud – Part 1

- Identity and Access Management
- Data Protection
- Lab 01 – Introduction to Security Policies

Module 4: Security IN the Cloud – Part 2

- Securing your infrastructure
- Monitoring and detective controls
- Lab 02 – Securing VPC Resources with Security Groups

Module 5: Security IN the Cloud – Part 3

- DDoS mitigation
- Incident response essentials
- Lab 03 – Automating Incident Response with AWS Config and AWS Lambda

Module 6: Course Wrap Up

- AWS Well-Architected tool overview

AWS Security Best Practices

Course description

Currently, the average cost of a security breach can be upwards of $4 million. AWS Security Best Practice provides an overview of some of the industry best practices for using AWS security and control types. This course helps you understand your responsibilities while providing valuable guidelines for how to keep your workload safe and secure. You will learn how to secure your network infrastructure using sound design options. You will also learn how you can harden your compute resources and manage them securely. Finally, by understanding AWS monitoring and

alerting, you can detect and alert on suspicious events to help you quickly begin the response process in the event of a potential compromise.

Level	Delivery method	Duration
Intermediate	Instructor-led training, presentations, demonstrations and assessments	1 day

Course objectives

In this course, you will learn to:

- Design and implement a secure network infrastructure
- Design and implement compute security
- Design and implement a logging solution

Intended audience

This course is intended for:

- Solution Architects, Cloud Engineers- Including Security Engineers, delivery and implementation engineers, professional services and Cloud Center of Excellence (CCOE)

Prerequisites

Students with a minimum of 1 years' experience managing open-source data frameworks, such as Apache Spark or Apache Hadoop will benefit from this course.

We recommend that attendees of this course have completed:

- AWS Security Essentials

AWS Security Best Practices Course Outline

Module 1: AWS Security Overview

- Shared Responsibility Model
- Customer Challenges
- Frameworks and Standards
- Establishing best practices
- Compliance in AWS

Module 2: Securing the Network

- Flexible and Secure
- Security inside the Amazon Virtual Private Cloud (Amazon VPC)
- Security Services
- Third party security Solutions

Lab 1: Controlling the Network

- Create a three-Security zone network infrastructure
- Implement network segmentation using security groups, Network Access Control Lists (NACLs) and public and private subnets
- Monitor network traffic to Amazon Elastic Compute Cloud (EC2) instances using VPC Flow Logs

Module 3: Amazon EC2 Security

- Compute Hardening

- Amazon Elastic Block Store (EBS) encryption
- Secure Management and maintenance
- Detecting vulnerabilities
- Using AWS Marketplace

Lab 2: Securing the Starting Point (EC2)

- Create a customer Amazon Machine Image (AMI)
- Deploy a new EC2 instance from a customer AMI
- Patch an EC2 instance using AWS Systems Manager
- Encrypt and EBS volume
- Understand how EBS encryption works and how it impacts other operations
- Use Security groups to limit traffic between EC2 instance to only that which is encrypted

Module 4: Monitoring and Alerting

- Logging Network Traffic
- Logging user and Application Programming Interface (API) traffic
- Visibility with Amazon CloudWatch
- Enhancing monitoring and alerting
- Verifying your AWS Environment

Lab 3: Security Monitoring

- Configuring and Amazon Linux 2 instance to send log files to Amazon CloudWatch
- Create Amazon CloudWatch alarms and notifications to monitor for failed login attempts
- Create Amazon CloudWatch alarms to monitor network traffic through a Network Address Translation (NAT) gateway

AWS Security Governance at Scale

Course description

Security is foundational to AWS. Governance at scale is a new concept for automating cloud governance that can help organization retire manual processes in account management, budget enforcement and security and compliance. By automating common challenges, organizations can scale without inhibiting agility, speed or innovation. In addition, they can provide decision makers with the visibility, control and governance necessary to protect sensitive date and systems.

In this course, you will learn how to facilitate developer speed and agility, and incorporate preventive and detective controls. By the end of this course, you will be able to apply governance best practices.

Level	Delivery method	Duration
Intermediate	Instructor-led training, presentations, demonstrations and assessments	1 day

Course objectives

In this course, you will learn to:

- Establish a landing zone with AWS Control Tower
- Configure AWS Organizations to create a multi-account environment
- Implement identity management using AWS Single Sign-On users and groups
- Federate access using AWS SSO
- Enforce polices using prepackaged guardrails
- Centralize logging using AWS CloudTrail and AWS Config
- Enable Cross-account security audits using AWS Identity and Access Management (IAM)
- Define workflows for provisioning accounts using AWS Service Catalog and AWS Security Hub

Intended audience

This course is intended for:

- Solution Architects, Security DevOps and Security Engineers

Prerequisites

We recommend that attendees of this course have:

- AWS Security Fundamentals
- Completed the AWS Security Essentials Course

AWS Security Governance at Scale Course Outline

Module 1: Governance at Scale

- Governance at Scale focus
- Business and Technical Challenges

Module 2: Governance Automation

- Multi-account strategies, guidance and architecture
- Environments for agility and governance at scale
- Governance with AWS Control Tower
- Use cases for governance at scale

Module 3: Preventive Controls

- Enterprise environment challenges for developers
- AWS Service Catalog
- Resource creation
- Workflows for provisioning accounts
- Preventive cost and security governance
- Self-Service with existing IT Service management (ITSM) Tools

Lab 1: Deploy Resources for AWS Catalog

- Create a new AWS Service Catalog portfolio and product
- Add an IAM roe to launch constraint to limit the actions the product can perform
- Grant access for an IAM role to view the catalog items
- Deploy an S3 bucket from an AWS Service Catalog Product

Module 4: Detective Controls

- Operations aspect of governance at scale
- Resource Monitoring
- Configuration rules for auditing
- Operational insights
- Remediation
- Clean up accounts

Lab 2: Compliance and Security Automation with AWS Config

- Apply Managed Rules through AWS Config to selected resources
- Automate remediation based on AWS Config rules
- Investigate the Amazon Config dashboard and verify resources and rule compliance

Lab 3: Taking Action with AWS Systems Manager

- Setup resource groups for various resources based on common requirements
- Perform automated actions against targeted resource groups

Module 5: Resources

- Explore additional resources for security governance at scale

Security Engineering on AWS

Course description

This course demonstrates how to efficiently use AWS security services to stay secure in the AWS Cloud. The course focuses on the security practices that AWS recommends for enhancing the security of your data and systems in the cloud. It highlights the security features of AWS key services including compute, storage, networking, and database services. You will also learn how to leverage AWS services and tools for automation, continuous monitoring and logging, and responding to security incidents.

Level	Delivery method	Duration
Intermediate	Instructor-led training, hands-on labs, and group exercises	3 days

Course objectives

In this course you will:

- Identify security benefits and responsibilities of using the AWS Cloud
- Build secure application infrastructures
- Protect applications and data from common security threats
- Perform and automate security checks
- Configure authentication and permissions for applications and resources
- Monitor AWS resources and respond to incidents
- Capture and process logs
- Create and configure automated and repeatable deployments with tools such as AMI's and AWS CloudFormation

Intended audience

This course is intended for:

- Security Engineers
- Security Architects
- Information Security professionals

Prerequisites

We recommend that attendees of this course have:

- Working knowledge of IT security practices and infrastructure concepts
- Familiarity with cloud computing concepts
- Completed AWS Security Essentials and Architecting on AWS Instructor Led Course

Security Engineering on AWS course outline

Day One

Module 1: Security on AWS

- Security in the AWS Cloud
- AWS Shared Responsibility Model
- Incident response overview
- DevOps with security engineering

Module 2: Identifying Entry Points on AWS

- Identify the different ways to access the AWS platform
- Understanding IAM policies
- IAM Permissions Boundary
- IAM Access Analyzer
- Multi-factor authentication
- AWS CloudTrail
- Hands-on lab 1: Cross-account access

Module 3: Security Considerations: Web Application Environments

- Threats in a three-tier architecture
- Common threats: User access
- Common threats: Data access
- AWS Trusted Advisor

Module 4: Application security

- Amazon machine images (AMIs)
- Amazon Inspector
- AWS Systems Manager
- Hands-on lab 2: Using AWS Systems Manager and Amazon Inspector

Module 5: Data security

- Data protection strategies
- Encryption on AWS
- Protecting data at rest with Amazon S3, Amazon RDS, and Amazon DynamoDB

- Protecting archived data with Amazon S3 Glacier
- Amazon S3 Analyzer
- Amazon S3 Access Points

Day Two

Module 6: Securing Network Communications

- Amazon VPC security considerations
- Amazon VPC traffic monitoring
- Responding to compromised instances
- Elastic Load Balancing
- AWS Certificate Manager

Module 7: Monitoring and Collecting Logs on AWS

- Amazon CloudWatch and CloudWatch Logs
- AWS Config
- Amazon Macie
- Amazon VPC Flow logs
- Amazon S3 server access logs
- ELB access logs
- Hands-on lab 3: Monitor and Respond with AWS Config

Module 8: Processing Logs on AWS

- Amazon Kinesis for log processing
- Amazon Athena for log processing
- Hands-on lab 4: Web Server Log Analysis

Module 9: Security Considerations: Hybrid Environments

- AWS Site-to-Site and Client VPN connections
- AWS Direct Connect
- AWS Transit Gateway

Module 10: Out-of-region protection

- Amazon Route 53
- AWS WAF
- Amazon CloudFront
- AWS Shield
- AWS Firewall Manager
- DDoS mitigation on AWS

Day Three

Module 11: Security Considerations: Serverless Environments

- Amazon Cognito
- Amazon API Gateway
- AWS Lambda

Module 12: Threat Detection and investigation

- Amazon GuardDuty
- AWS Security Hub

- AWS Detective

Module 13: Secrets Management on AWS

- AWS Key Management Service (AWS KMS)
- AWS CloudHSM
- AWS Secrets Manager
- Hands-on lab 5: Using AWS KMS

Module 14: Automation and Security by Design

- AWS CloudFormation
- AWS Service Catalog
- Hands-on lab 6: Security automation on AWS with AWS Service Catalog

Module 15: Account Management and Provisioning on AWS

- AWS Organizations
- AWS Control Tower
- AWS SSO
- AWS Directory Services
- Hands-on lab 7: Federated Access with ADFS

Exam Readiness

Before taking an AWS Certification exam, we recommend you have hands-on experience with relevant AWS products and services. We offer a number of resources to supplement your experience and help you prepare for your AWS Certification. For Exam Preparation Resources please visit www.skillbuilder.aws