



G-Cloud 15

Amazon Web Services EMEA SARL, UK Branch
(AWS)

Service Description

Lot 3 Cloud Support Services

AWS Training



Table of Contents

1	How to Use the AWS Service Definition Documents	1
2	Accessibility	1
3	Digital Training: AWS Skill Builder	2
4	AWS Classroom Training – High-level Overview	3
5	AWS Foundational Courses	6
5.1	AWS Technical Essentials	6
5.2	AWS Cloud Practitioner Essentials	8
5.3	AWS Cloud Essentials for Business Leaders	11
6	AWS Technical Courses – Grow Your Skills	12
6.1	Architect Course Curriculum	12
6.2	Artificial Intelligence Course Curriculum	23
6.3	Containers Course Curriculum	46
6.4	Data Analytics Course Curriculum	48
6.5	Databases Course Curriculum	62
6.6	Developer Course Curriculum	64
6.7	DevOps Course Curriculum	69
6.8	Migrate and Transfer Course Curriculum	72
6.9	Networking and Content Delivery Course Curriculum	77
6.10	Operations Course Curriculum	79
6.11	Serverless Course Curriculum	82
6.12	Security Course Curriculum	85
6.13	Storage Course Curriculum	89
7	Exam Readiness	92

This document is provided for informational purposes only. This document contains Amazon Confidential Information and is subject to the terms of the parties' agreement governing confidentiality, which include the terms governing Confidential Information in the AWS Customer Agreement (aws.amazon.com/agreement) or other agreement between you and AWS ("Service Agreement") governing Confidential Information. Where the parties have not entered into a Service Agreement, the terms of the AWS Customer Agreement control the Confidential Information provided in this document. Unless required by law, you agree not to disclose the contents of this document. Where you are required to disclose document contents by law, you will provide AWS with written notice and an opportunity to seek redactions or otherwise prevent disclosure of the information to the maximum extent permitted by law. The offerings described in this document represent AWS's current products and practices as of the date of issue of this document and are subject to change. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers. For current prices for AWS services, please refer to the AWS website at www.aws.amazon.com.

1 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:

- IaaS and PaaS (Lot 1a); SaaS and SaaS (Lot 2a); SaaS (Lot 2b)
- Professional Services (Lot 3)
- Support Services (Lot 3)
- **Training Services (Lot 3)**
- AWS Managed Services (Lot 3)

This **AWS Training Service Description** document describes the key features for each of the AWS Training offerings available to Customers on G-Cloud 15 in Lot 3.

To access the options 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 service as an option that can be procured under their G-Cloud 15 Call-Off Contract.

Note: Lot 2a (Infrastructure Software as a Service - I-SaaS), Lot 2b (Software-as-a-Service) and Lot 3 (Cloud Support Services) services may be purchased as ancillary services to Lot 1 (IaaS and PaaS) offerings. The framework's Full Taxonomy of Services uses a Lot availability key where services are marked with "O" (Optional) to indicate they "may be purchased as ancillary service to Core Requirements".

This means that while Lot 2a, Lot 2b, and Lot 3 services are primarily available through their respective dedicated lots, buyers have the flexibility to procure these services alongside their Lot 1a or Lot 1b cloud hosting purchases when they complement or support the core infrastructure and platform services. This ancillary purchasing mechanism allows buyers to create comprehensive cloud solutions by bundling infrastructure software and support services with their primary IaaS/PaaS deployments, streamlining procurement and ensuring integrated service delivery from a single supplier when appropriate.

AWS recommends that Buyers list all of the Digital Marketplace Service ID's for every service across all relevant lots in its Call-Off Contract to enable the option to switch between Services flexibly during the term. For a list of all AWS Digital Marketplace Service ID's, please contact an AWS account representative at aws-ukps-frameworks@amazon.com.

Please note that the options or parameters selected by AWS on this framework are those that most closely align with our existing commercial services. AWS is happy to provide additional information about our services upon request.

To find out more about AWS on G-Cloud and AWS Cloud services, visit us at [AWS on G-Cloud UK](#).

2 Accessibility

All AWS Training and Certification Course Materials meet the WCAG 2.1 AA Accessibility Standard.

3 Digital Training: AWS Skill Builder

AWS Skill Builder provides self-paced digital training on AWS services and is available any time you need it.

[AWS Skill Builder Free Tier](#): AWS Skill Builder Free Tier lets your employees explore AWS learning plans and 600+ digital courses over 30 technical domains. This can help them own their career and achieve goals on schedule.

[AWS Skill Builder Team Subscription](#): AWS Skill Builder Team subscription is a digital training subscription for organisations. The Team subscription provides access to AWS self-paced digital training content including digital classrooms, self-paced labs, game-based learning, and AWS Certification Exam preparation. It also provides administrative capabilities to assign and track training for individuals and teams.

Why should my organisation use AWS Skill Builder Subscription?

With AWS Skill Builder Team subscription, individuals and teams gain access to a broad portfolio of learning experiences, all assignable and trackable from a self-service dashboard. Subscribers gain hands-on experience with real AWS services through access to AWS Builder Labs. They can learn role-based skills through an immersive game, AWS Cloud Quest, set in a virtual world. Teams can challenge themselves to solve practical problems with AWS Jam. Organisations looking to grow their ranks of certified individuals gain the benefits of exam preparation courses and practice exams for AWS Certification. These features help you upskill your workforce faster and track their progress in a self-service manner.

How many training resources are included with AWS Skill Builder Team subscription?

Learners who have the Team subscription have access to over 600 self-paced digital courses that are freely available to all AWS Skill Builder learners. In addition to these courses, Team subscription includes:

AWS Digital Classrooms. The same knowledge and skills as live classroom training through recorded expert instructors, demonstrations, knowledge checks, assessments, and hands-on labs for a near-classroom experience in a flexible self-paced format. 16 different classes are currently available.

200+ guided AWS Builder Labs. These build practical skills within the AWS Management Console Environment, including new labs covering generative AI, machine learning, data analytics, and security.

AWS Learning Assistant. An AI-powered chat guide providing expert guidance for over 200 self-paced labs, offering real-time support and personalized learning assistance.

45+ learning plans to accelerate skills development. These provide a curated list of courses and labs related to a job-role or a domain area. Learners can enrol in learning plans and complete assessments to earn digital badges that can be shared on social media. These badges demonstrate skills and knowledge for specific AWS Cloud topics

AWS Cloud Quest. Learners can develop AWS skills in a fun environment, where they can zap drones, befriend pets, and solve puzzles. The course helps users learn real solutions inside a live AWS environment, and develop skills through role-based immersion, covering 7+ cloud job roles and 75+ solutions quests

AWS Card Clash. This course lets you test your Cloud Architecture Knowledge in an immersive virtual card game. Now includes multi-player features for competitive learning.



AWS Escape Room. Learners can prepare for the Cloud Practitioner exam in 3D virtual escape room adventure

AWS Industry Quest. Build solutions for common use-cases in specific industries, such as Financial Services, Healthcare, and Manufacturing.

AWS SimuLearn. AI-powered meeting simulator for hands-on practice in realistic business scenarios.

AWS Skill Builder Trivia. Interactive learning experience to test and reinforce AWS knowledge.

AWS Jam. Offers 700+ scenario-based challenges that enable teams and cohorts to practice and validate skills. Users can create their own Jam Journeys and invite colleagues to compete in Jam Tournaments.

Enhanced Exam Preparation Courses. Curated content, practice questions, and builder labs to prepare for certification, including preparation for the new AWS Certified Generative AI Developer – Professional certification.

AWS Certification Official Practice Exams. Full-length practice exams with same style, rigor and scoring as AWS certification style exams, now available in multiple languages.

Administrator Dashboard & Reporting. Assign training to users and groups, monitor progress, and track training with built-in comprehensive reporting. Enhanced with organisation and team-level analytics, automated seat assignment, and bulk user import via CSV.

Single Sign-On Support (Optional, for 30+ seats). Supports SAML2.0 protocol using the following identity providers: Microsoft ADFS, Azure AD, Ping and Okta.

4 AWS Classroom Training – High-level Overview

With AWS Classroom Training, your learners can take in-person and virtual training from instructors who teach 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 Cloud Practitioner Essentials	A fundamental course for individuals seeking overall understanding of AWS Cloud regardless of technical role. Provides overview of cloud concepts, AWS services, security, architecture, pricing, and support.	1 day
AWS Fundamentals	AWS Technical Essentials	Learn fundamentals to identify AWS services and make informed decisions about IT solutions based on business requirements. Get started working on AWS.	1 day
AWS Fundamentals	AWS Cloud Essentials for Business Leaders	Learn fundamental cloud computing concepts and how cloud strategy can help meet business objectives. Explores advantages and addresses concepts like security and compliance.	0.5 day



Skill Domain	Course Title	Description	Duration
Architecting	Architecting on AWS	Learn to build IT infrastructure on AWS and optimize the AWS Cloud by understanding AWS services and how they fit into cloud-based solutions.	3 days
Architecting	AWS Well-Architected Best Practices	Learn the AWS Well-Architected Framework and its 5 pillars, covering the Well-Architected Review Process using the AWS Well-Architected Tool.	1 day
Architecting	Advanced Architecting on AWS	Building on Architecting on AWS, learn to build complex solutions incorporating data services, governance, and security on AWS.	3 days
Architecting	AWS Advanced Well-Architected Best Practices	Learn phases of a review, including preparation, execution, and guidance after review completion. For those familiar with intermediate AWS Well-Architected Best Practices.	1 day
Artificial Intelligence	Generative AI Essentials on AWS	Learn fundamental concepts, methods, and strategies for using generative AI and addressing business needs.	1 day
Artificial Intelligence	Generative AI for Executives	Learn to leverage generative AI within your organisation, covering business value, use cases, and implementation considerations.	0.5 days
Artificial Intelligence	Agentic AI Foundations	Explore core principles for designing Agentic AI systems using AWS services. Learn how Agentic AI differs from traditional systems and build autonomous solutions.	1 day
Artificial Intelligence	Developing Generative AI Applications on AWS	Explore core principles for designing Agentic AI systems using AWS services. Learn how Agentic AI differs from traditional systems and build autonomous solutions.	2 days
Artificial Intelligence	Advanced Generative AI Development on AWS	Build expertise across the entire generative AI stack from foundation models to enterprise integration patterns.	3 days
Artificial Intelligence	Machine Learning Engineering on AWS	Learn to build, deploy, orchestrate, and operationalize ML solutions at scale using AWS services like Amazon SageMaker and Amazon EMR through theory, practical labs, and hands-on activities.	3 days
Artificial Intelligence	MLOps Engineering on AWS	Learn the importance of data, model and code for successful ML deployments. Demonstrates tools, automation, processes and teamwork for ML operations.	3 days
Artificial Intelligence	Practical Data Science with Amazon SageMaker	Learn to solve real-world use cases with Machine Learning and produce actionable results using Amazon SageMaker.	1 day



Skill Domain	Course Title	Description	Duration
Artificial Intelligence	Amazon SageMaker Studio for Data Scientists	Learn how SageMaker Studio helps data scientists prepare, build, train, deploy, and monitor ML models by bringing together capabilities purpose-built for ML to improve productivity throughout the ML lifecycle.	3 days
Containers	Running Containers on Amazon Elastic Kubernetes Services (EKS)	Learn container management and orchestration for Kubernetes using Amazon EKS.	3 days
Data Analytics	Building Data Lakes on AWS	Learn to build operational data lakes supporting analysis of structured and unstructured data, including components and functionality of relevant services.	1 day
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.	1 day
Data Analytics	Building Data Analytic Solutions using Amazon Redshift	Learn to integrate Amazon Redshift with data lakes to support analytics and machine learning workloads.	1 day
Data Analytics	Building Streaming Data Analytics Solutions on AWS	Learn to build streaming data analytics solutions using AWS services like Amazon Kinesis and Amazon MSK.	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	Data Engineering on AWS	Learn to design, build, optimize, and secure data engineering solutions using AWS services.	3 days
Databases	Building Modern Applications with AWS NoSQL Databases	Use AWS purpose-built databases to build modern applications 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
DevOps	DevOps Engineering on AWS	Learn common DevOps patterns to develop, deploy, and maintain applications on AWS.	3 days
Migrate and Transfer	AWS Migration Essentials	Learn foundational knowledge for planning and performing lift-and-shift migrations to AWS Cloud.	1 day
Migrate and Transfer	Migrating to AWS	Learn cloud migration strategies and how to apply each migration process step.	3 days

Skill Domain	Course Title	Description	Duration
Networking and Content Delivery	Networking Essentials for Cloud Applications on AWS	Learn networking concepts and services within AWS cloud environment 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 fundamental AWS Cloud security concepts, including access control, data encryption, and network security.	1 day
Security	Security Engineering on AWS	Learn how to efficiently use AWS security services to stay secure in the AWS Cloud.	3 days
Storage	Designing and Implementing Storage on AWS	Learn where, how, and when to use different AWS storage services and solutions for storing, accessing, and protecting data.	3 days

5 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.

5.1 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.

- Course level: Fundamental
- Duration: 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 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 (Amazon 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 to customers
- Individuals interested in learning how to get started with AWS
- SysOps administrators
- Solutions architects
- Developers

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 is required.

AWS Technical Essentials Course Outline

Module 1: AWS 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

5.2 AWS Cloud Practitioner Essentials

Course Description

This course is for individuals who seek an overall understanding of the Amazon Web Services (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. This course also helps you prepare for the AWS Certified Cloud Practitioner exam.

- Course level: Fundamental
- Duration: 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 organisation'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 Organisations
- 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 Organisations 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

5.3 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.

- Course level: Fundamental
- Duration: 1 day

Course Objectives

This course is you will learn to:

- Explain the role of information technology (IT) in an organisation 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

Introduction to Cloud Essentials for Business Leaders Course

Module 2: Information Technology for Business Transformation

- Role of IT in an organisation 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.

6 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

6.1 Architect Course Curriculum

6.1.1 Architecting on AWS

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.

- Course level: Intermediate
- Duration: 3 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
- 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:

- 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 Course Outline

Day 1

Module 1: Architecting Fundamentals

- AWS services
- AWS infrastructure
- AWS Well-Architected Framework
- Hands-on lab: Explore and interact with the AWS Management Console and AWS Command Line Interface

Module 2: Account Security

- Principles and identities
- Security policies
- Managing multiple accounts

Module 3: Networking 1

- IP addressing
- VPC fundamentals
- VPC traffic security

Module 4: Compute

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



Day 2

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: Create a database layer in your Amazon VPC infrastructure

Module 7: Monitoring and Scaling

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

Module 8: Automation

- AWS CloudFormation
- Infrastructure management

Module 9: Containers

- Microservices
- Containers
- Container services

Day 3

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
- Hands-on Lab: Build a serverless architecture

Module 12: Edge Services

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

Module 13: Backup and Recovery

- Disaster planning
- AWS Backup
- Recovery strategies
- Hands-on Lab: Capstone lab – 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.

6.1.2 AWS Well-Architected Best Practices

Course Description

The AWS Well-Architected Framework helps you make informed decisions about your customers' architectures cloud-focused way and understand the impact of design decisions. By using the Well-Architected Framework, you will understand the risks in your architecture and ways to mitigate them. This course provides a deep dive into the AWS Well-Architected Framework and its six pillars. This course also covers the Well-Architected review process and using the AWS Well-Architected Tool to complete reviews.

- Course level: Intermediate
- Duration: 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

- Brief history of AWS Well-Architected
- AWS Well-Architected pillars
- Design principles
- Applying the AWS Well-Architected Framework
- AWS Well-Architected Tool

Module 2: Operational Excellence

- Operational Excellence design principles
- Case study
- Hands-On Lab: Operational Excellence

Module 3: Reliability

- Reliability design principles
- Hands-On Lab: Reliability

Module 4: Security

- Security design principles
- Hands-On Lab: Security

Module 5: Performance Efficiency

- Performance Efficiency design principles
- Hands-On Lab: Performance Efficiency

Module 6: Cost Optimization

- Cost Optimization design principles
- Hands-On Lab: Cost Optimization

Module 7: Sustainability

- Sustainability design principles
- Sustainability best practices
- Sustainability pillar resources

Module 8: Course Summary

- Recap
- Resources



- Continue your learning

6.1.3 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.

- Course level: Advanced
- Duration: 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 Organisations organisational 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 5G networks, 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 Organisations 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

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

6.1.4 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.

- Course level: Advanced
- Duration: 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

6.2 Artificial Intelligence Course Curriculum

6.2.1 Generative AI Essentials on AWS

Course Description

In this course, you will learn about the fundamental concepts, methods, and strategies for using generative AI. You will gain a solid understanding of use cases where generative AI can provide solutions and address business needs. Additionally, you will learn about practical insights into technologies related to generative AI and how you can use those technologies to solve real-world problems. By the end of the course, you will explore project planning and how to discuss implementation of generative AI in your organisation.

- Course level: Fundamental
- Duration: 1 day

Course Objectives

In this course, you will learn to:

- Summarize generative AI concepts, methods, and strategies
- Discuss the appropriate use of generative AI and machine learning and their technologies

- Describe how to use generative AI responsibly and safely
- Recognize the types of generative AI solutions with specific use cases
- Explain implementation and project planning of generative AI to your organisation

Intended Audience

This course is intended for those with limited prior knowledge of Generative AI.

Prerequisites

- None

Generative AI Essentials on AWS Course Outline

Module 1: Introducing Generative AI

- Generative AI explained
- Foundation models
- AWS generative AI services
- Demo: Generative AI solution

Module 2: Exploring Generative AI Use Cases

- Identify suitable use cases
- Generative AI applications and use cases
- Explore generative AI use case scenarios
- Use case for class

Module 3: Essentials of Prompt Engineering

- Introduction to prompt engineering
- Prompt design best practices
- Advanced prompting strategies
- Model settings and parameters
- Hands-on Lab: Optimizing Slogan Generation with Amazon Bedrock

Module 4: Responsible AI Principles and Considerations

- Introduction to responsible AI
- Core dimensions of responsible AI
- Generative AI considerations
- Hands-on Lab: Implementing Responsible AI Principles with Amazon Bedrock Guardrails

Module 5: Security, Governance, and Compliance

- Security overview
- Adverse prompts
- Generative AI security services
- Governance
- Compliance

Module 6: Implementing Generative AI Projects

- Introduction – Generative AI application
- Define a use case
- Select a foundational model
- Improve performance
- Evaluate results
- Deploy the application
- Demo: Amazon Q Business

Module 7: Integrating Generative AI into the Development Lifecycle Introduction

- Hands-on Lab: Capstone – Creating a Project Plan with Generative AI

Module 8: Course Wrap Up

6.2.2 Generative AI for Executives

Course Description

In this course, you will learn how to leverage generative artificial intelligence (AI) within your organisation. We'll cover how to drive business value with generative AI, the use cases across various industries, and the considerations to implement generative AI safely and responsibly. The goal of this course is to provide you with the fundamental concepts and tools you'll need to successfully lead generative AI initiatives within your organisation.

- Course level: Fundamental
- Duration: 0.5 Day

Course Objectives

In this course, you will learn to:

- Recognize the potential business value of generative AI
- Identify real world use cases that you can implement today
- Manage the people, process, and technology changes needed to be successful
- Use generative AI safely and responsibly
- Identify the specific steps you can take to get started with generative AI

Intended Audience

This course is intended for Executives and Senior Business Leaders.

Prerequisites

- None

Generative AI for Executives Course Outline

Module 1: Introduction to Generative AI

- Definitions and terminology
- AWS approach to generative AI



Module 2: Generative AI Use Cases

- Common use cases
- Real-world case studies

Module 3: Overcoming Technical and Organisational Challenges

- Security
- Accuracy
- Cost
- People and culture

Module 4: Implementation

- Identifying your use case
- Assessing data, technology, people, and processes
- Evaluating business impact and scaling

Module 5: Next Steps in Your Generative AI Journey

- Next steps and additional resources
- Course summary

6.2.3 Agentic AI Foundations

Course Description

In this course, you'll explore the core principles and strategies for designing Agentic AI systems using AWS services. You'll learn how Agentic AI differs from traditional conversational systems, and how to use tools like Amazon Q, Kiro, Amazon Bedrock Agents, and Amazon Bedrock AgentCore to build autonomous, goal-driven solutions that solve real-world problems.

- Course level: Fundamental
- Duration: 1 day

Course Objectives

In this course, you will learn to:

- Summarize the evolution of Agentic AI and define what makes something "agentic"
- Identify core components of agentic systems
- Distinguish between workflow, autonomous, and hybrid agents
- Compare AWS service options for Agentic AI
- Describe capabilities and use cases of Amazon Q Developer, Amazon Q Business, and Kiro
- Explain Amazon Bedrock AgentCore and Amazon Bedrock Agents fundamentals
- Identify basic implementation patterns for Agentic AI
- Describe observability and interoperability patterns for production agentic AI systems

Intended Audience

This course is intended for:

- Software developers new to Agentic AI seeking foundational knowledge



- Technical professionals exploring AI capabilities and interested in core components and applications of agentic AI
- Development teams evaluating Agentic AI solutions and needing to differentiate between agent types
- AWS Users expanding into Agentic AI, including current users of Amazon Q Developer, Amazon Q Business, and Amazon Bedrock Agents

Prerequisites

We recommend that attendees of this course have:

- Generative AI Essentials or equivalent work experience
- Basic AWS knowledge and software development experience

Agentic AI Foundations Course Outline

Module 1: From LLMs to Agents

- Understanding Large Language Models (LLMs)
- Innovations powering agents
- Evolution timeline from LLMs to Agents

Module 2: Exploring Agentic AI

- Understanding Agentic AI
- Types of AI agents
- Agentic AI applications

Module 3: Understanding Agentic AI Workflows

- Workflow patterns
- Amazon Bedrock flows overview

Module 4: Introducing Autonomous Agents

- How Autonomous Agents work
- ReAct
- ReWoo
- Multi-agent collaboration
- AWS Agentic AI solutions

Module 5: Amazon Q and Agentic Development Tools

- Amazon Q Developer
- Amazon Q Business
- Amazon Q in AWS Services
- Kiro: AI-powered IDE with spec-driven development

Module 6: Agentic AI with Amazon Bedrock

- Amazon Bedrock Agents
- Amazon Bedrock AgentCore
- Hands-on lab: Explore Amazon Bedrock Agents integrated with Amazon Bedrock Knowledge Bases and Amazon Bedrock Guardrails

Module 7: Building DIY Solutions

- DIY solutions
- Observability and Monitoring
- Agent Interoperability

Module 8: Course Wrap Up

6.2.4 Building Agentic AI with Amazon Bedrock Agent Core

Course Description

In this course, you'll explore the core principles and strategies for designing Agentic AI systems using AWS services. You'll learn how Agentic AI differs from traditional conversational systems, and how to use tools like Strands Agents SDK, and Amazon Bedrock AgentCore to build autonomous, goal-driven solutions that solve real-world problems.

- Course level: Fundamental
- Duration: 1 day

Course Objectives

In this course, you will learn to:

- Define agentic AI characteristics and differentiate them from traditional AI systems.
- Identify the core agent components and their interactions.
- Describe how Bedrock AgentCore services support agentic AI.
- Deploy agents by using supported frameworks with AgentCore Runtime.
- Describe the core features of AgentCore Runtime.
- Configure serverless execution with session isolation.
- Configure AgentCore Identity for enterprise security requirements.
- Create policies to secure agent tool calls using AgentCore Policy.
- Implement secure token management and permission delegation.
- Ensure compliance with data governance and audit requirements.
- Implement different tool integration patterns, including built-in tools and protocol-based tools.
- Design and deploy Model Context Protocol (MCP) servers and clients for extensible agent capabilities.
- Describe common authentication patterns for agent tool use.
- Configure AgentCore Gateway components for secure and authorized tool access.
- Implement agentic memory patterns for different use cases.
- Configure AgentCore Memory operations for context-aware development.
- Optimize memory performance for production workloads.
- Configure AgentCore Observability for production monitoring.
- Implement Amazon CloudWatch integration and specialized tracing.
- Describe the core features of AgentCore Evaluations.
- Integrate agentic systems with production APIs and services.
- Design deployment strategies for production environments.

- Assess production readiness and establish continuous improvement processes

Intended Audience

This course is intended for:

- Software developers seeking intermediate knowledge for building Agentic AI systems.
- Technical professionals exploring AI capabilities and interested in building Agentic AI systems.
- Development teams building Agentic AI solutions.

Prerequisites

We recommend that attendees of this course have:

- Agentic AI Foundations

Building Agentic AI with Amazon Bedrock Agent Core Course Outline

Module 1: Foundations of Agentic AI Patterns

- Agent building blocks
- Amazon Bedrock AgentCore introduction

Module 2: AgentCore Runtime and Framework Integration

- Supported frameworks and implementation
- AgentCore Runtime overview
- Infrastructure and deployment

Module 3: Security and Identity Management

- Security and identity management
- Securing your agents with AgentCore Identity

Module 4: Tool Integration and AgentCore Gateway

- Amazon Bedrock AgentCore Policy
- Built-in tools and custom integration
- Model Context Protocol (MCP)
- AgentCore Gateway
- Implementing AgentCore Gateway
- Amazon Bedrock AgentCore Policy

Module 5: Agentic Memory Implementation

- Agentic memory core concepts
- AgentCore Memory
- Securing AgentCore Memory
- Hands-on Lab: Enhance and Scale Agents with Amazon Bedrock AgentCore

Module 6: Production Monitoring and Observability

- Monitoring agents with AgentCore Observability
- Verifying agent performance with AgentCore Evaluation

Module 7: Course Wrap-Up

6.2.5 Building Advanced Agentic Systems on AWS

Course Description

In this course, you'll learn about implementing production-ready multi-agent systems using Amazon Bedrock AgentCore, covering multi-agent patterns, context optimization techniques, security configurations, and monitoring frameworks. You will examine the skills needed to move beyond proof-of-concept to scalable, secure, and observable agentic AI implementations. The course prepares you to design and deploy advanced agentic systems ready for real-world production environment.

- Course level: Advanced
- Duration: 1 day

Course Objectives

In this course, you will learn to:

- Analyze scenarios that require multi-agent architectures.
- Describe primary multi-agent communication patterns and their use cases.
- Configure agent-as-tool patterns for production deployments.
- Implement memory sharing using available platform capabilities.
- Implement context management strategies for production workloads.
- Design context compression and optimization mechanisms.
- Optimize resource usage and cost management across multi-agent systems.
- Configure policy-based access control using AgentCore Policy Engine.
- Implement VPC integration for secure agent deployments.
- Implement distributed tracing and monitoring across multi-agent systems.
- Establish comprehensive agent evaluation frameworks.
- Configure integration patterns for enterprise observability systems
- Establish comprehensive audit trails and compliance monitoring.
- Integrate agentic systems with production APIs and services.
- Design deployment strategies for production environments.
- Assess production readiness and establish continuous improvement processes.

Intended Audience

This course is intended for:

- Software developers seeking intermediate knowledge for building advanced agentic AI systems
- Technical professionals exploring AI capabilities and interested in building advanced agentic AI systems



- Development teams building advanced agentic AI systems

Prerequisites

We recommend that attendees of this course have:

- Agentic AI Foundations
- Building Agentic AI with Amazon Bedrock AgentCore

Building Advanced Agentic Systems on AWS Course Outline

Module 1: Multi-Agent Architecture and Communications Patterns

- Single agent limitations and multi-agent benefits
- Instructor Demonstration: Building and Deploying Intelligent Financial Agents with Amazon Bedrock Strands and AgentCore - Task 1: Building a Personal Budget Assistant with Strands Agents
- Multi-agent communication patterns
- Memory and state management
- Instructor Demonstration: Building and Deploying Intelligent Financial Agents with Amazon Bedrock Strands and AgentCore - Task 2: Building a Multi-Agent System for Complex Financial Analysis
- Instructor Demonstration: Building and Deploying Intelligent Financial Agents with Amazon Bedrock Strands and AgentCore

Module 2: Context Engineering and Performance Optimization

- Context as finite resource
- Context optimization techniques
- Tool design for context efficiency

Module 3: Security and Compliance Implementation

- Production Identity and Access Management
- VPC integration and network security

Module 4: Production Monitoring, Observability, and Evaluation

- Monitoring architecture
- AgentCore evaluation
- Enterprise observability integration
- Instructor Demonstration: Building and Deploying Intelligent Financial Agents with Amazon Bedrock Strands and AgentCore - Task 3: Deploying Production-Ready Agents with Amazon Bedrock AgentCore

Module 5: Well-Architected Agentic AI Systems

- Applying the Well-Architected framework
- Well-Architected deployment
- Production readiness

Module 6: Course Wrap-up



6.2.6 Developing Generative AI Applications on AWS

Course Description

In this advanced two-day course, software developers learn to build and customize AI solutions by using Amazon Bedrock programmatically. Through hands-on exercises and labs, participants will invoke foundation models through Amazon Bedrock APIs, implement Retrieval Augmented Generation (RAG) patterns with Amazon Bedrock Knowledge Bases, and develop AI agents with tool integration. The course focuses on the practical implementation of prompt engineering techniques, responsible AI practices with Amazon Bedrock Guardrails, open-source framework integration, and architectural patterns for real-world business applications.

- Course level: Advanced
- Duration: 2 days

Course Objectives

In this course you will learn to:

- Develop generative AI applications using Amazon Bedrock.
- Design architecture patterns for generative AI applications.
- Configure Amazon Bedrock APIs to invoke foundation models (FMs) programmatically. Develop agentic AI applications by integrating Amazon Bedrock tools and open-source frameworks.
- Build custom solutions with Retrieval Augmented Generation (RAG) and Amazon Bedrock Knowledge Bases.
- Integrate open-source SDKs with Amazon Bedrock to build business.
- Optimize model responses by applying prompt engineering techniques.
- Evaluate generative AI application components.
- Implement responsible AI practices to protect generative AI.
- 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.

Prerequisites

We recommend that attendees of this course have:

- Completed the Generative AI Essentials AWS instructor-led course
- Intermediate-level proficiency in Python



- Familiarity with AWS Cloud

Developing Generative AI Applications on AWS Course Outline

Day 1

Course Introduction

Module 1: Exploring Components of Generative AI Applications on AWS

- Understanding generative AI concepts
- Identifying AWS generative AI stack components
- Designing generative AI application components

Module 2: Programming with Amazon Bedrock

- Guiding model response generation
- Using Amazon Bedrock programmatically
- Hands-on lab: Develop with Amazon Bedrock APIs
- Hands-on lab: Develop Streaming Patterns with Amazon Bedrock APIs

Module 3: Applying Prompt Engineering for Developers

- Introducing prompt engineering
- Introducing prompt techniques
- Optimizing prompts for better results

Module 4: Using Amazon Bedrock APIs in Common Architectures

- Implementing architecture patterns with Amazon Bedrock APIs
- Exploring common use cases
- Adding conversational memory to extend context
- Hands-on lab: Develop Conversation Patterns with Amazon Bedrock APIs

Module 5: Customizing Generative AI Responses with RAG

- Implementing Retrieval Augmented Generation (RAG)
- Using Amazon Bedrock Knowledge Bases
- Hands-on lab: Develop Retrieval Augmented Generation (RAG) Applications with Amazon Bedrock Knowledge Bases

Module 6: Integrating Open-Source Frameworks with Amazon Bedrock

- Invoking a foundation model in Amazon Bedrock using LangChain
- Using LangChain for context-aware responses
- Hands-on lab: Develop a Generative AI Application Pattern using Open-Source Frameworks and Amazon Bedrock Knowledge Bases

Day 2

Module 7: Evaluating Generative AI Application Components

- Evaluating application components
- Evaluating model output
- Evaluating RAG output

- Optimizing latency and cost
- Hands-on lab: Evaluating Retrieval Augmented Generation (RAG) Applications

Module 8: Implementing Responsible AI

- Understanding responsible AI
- Mitigating bias and addressing prompt misuses
- Using Amazon Bedrock Guardrails
- Hands-on lab: Securing Generative AI Applications Using Bedrock Guardrails

Module 9: Using Tools and Agents in Generative AI Applications

- Using tools
- Understanding AI agents
- Understanding open-source agentic frameworks
- Understanding agent interoperability

Module 10: Developing Amazon Bedrock Agents

- Implementing Amazon Bedrock Flows
- Designing Amazon Bedrock Agents
- Developing Amazon Bedrock Inline Agents
- Designing multi-agent collaboration
- Using Amazon Bedrock AgentCore
- Hands-on lab: Developing Amazon Bedrock Agents Integrated with Amazon Bedrock Knowledge Bases and Guardrails

Module 11: Course Wrap Up

6.2.7 Advanced Generative AI Development on AWS

Course Description

The Advanced Generative AI Development on AWS is designed for developers seeking to master the implementation of production-ready generative AI solutions on AWS. The course addresses the needs of organisations embarking on their generative AI journey and how to build comprehensive generative AI strategies that align with broader business objectives.

This advanced 3-day instructor-led training builds expertise across the entire generative AI stack - from foundation models to enterprise integration patterns. In addition, you will learn about advanced data processing techniques, vector database implementation and retrieval augmentation, sophisticated prompt engineering and governance, agentic AI systems and tool integration, AI safety and security measures, performance optimization and cost management strategies, comprehensive monitoring and observability solutions, testing and validation frameworks.

The course structure follows AWS's proven model for generative AI adoption, progressing from experimentation to production-ready implementations.

- Course level: Advanced
- Duration: 3 days

Course Objectives

In this course, you will learn to:

- Develop production-ready generative AI solutions using AWS services that meet enterprise requirements for security, scalability, and reliability
- Evaluate and select appropriate foundation models for specific business use cases, including benchmarking performance and implementing dynamic model selection architectures
- Design and implement foundation model systems with circuit breakers, cross-region deployment, and degradation strategies
- Build comprehensive data processing pipelines for multi-modal inputs, including validation workflows and optimization techniques
- Implement sophisticated vector database solutions using Amazon Bedrock Knowledge Bases, OpenSearch, and hybrid approaches for effective retrieval augmentation
- Create and manage advanced prompt engineering frameworks, including chain-of-thought reasoning and enterprise-wide prompt governance systems
- Explain components of Agentic AI frameworks and Amazon Bedrock AgentCore
- Implement comprehensive AI safety and security controls, including content filtering, privacy preservation, and adversarial testing mechanisms
- Optimize performance and manage costs through token efficiency strategies, batching implementations, and intelligent caching systems
- Design and implement comprehensive monitoring and observability solutions for foundation model applications
- Create systematic testing and validation frameworks for continuous quality assurance of AI applications
- Integrate generative AI solutions within enterprise environments using secure, compliant, and scalable architectural patterns

Intended Audience

This course is intended for:

- Software Developers
- Technical Professionals

Prerequisites

We recommend that attendees of this course have:

- AWS Technical Essentials
- Generative AI Essentials on AWS
- 2 or more years of experience building production grade applications on AWS or with open-source technologies, general AI/ML or data engineering experience
- 1 year of hands-on experience implementing generative AI solutions

Advanced Generative AI Development on AWS Course Outline

Day 1

Module 1: Foundation Model Selection and Configuration

- Enterprise foundation model evaluation framework



- Dynamic model selection architecture patterns
- Resilient foundation model system designs
- Cost optimization and economic modelling

Module 2: Advanced Data Processing for Foundation Models

- Comprehensive data validation and quality assurance
- Multi-modal data processing pipelines
- Input optimization and performance enhancement

Module 3: Vector Databases and Retrieval Augmentation

- Enterprise vector database architecture
- Advanced document processing and chunking strategies
- Sophisticated retrieval system implementation
- Hands-on Lab: Develop Retrieval Augmented Generation (RAG) Applications with Amazon Bedrock Knowledge Bases

Day 2

Module 4: Prompt Engineering and Governance

- Advanced prompt engineering frameworks
- Complex prompt orchestration systems
- Enterprise prompt governance and management
- Hands-on Lab: Develop conversation pattern with Amazon Bedrock APIs

Module 5: Implementing Agentic AI Frameworks with Amazon Bedrock AgentCore

- Agentic AI Frameworks
- Amazon Bedrock AgentCore

Module 6: AI Safety and Security

- Comprehensive content safety implementation
- Privacy-preserving AI architecture
- AI governance and compliance frameworks

Day 3

Module 7: Performance Optimization and Cost Management

- Token efficiency and cost optimization
- High-performance system architecture
- Intelligent caching systems implementation
- Hands-on Lab: Building Secure and Responsible Gen AI with Guardrails for Amazon Bedrock

Module 8: Monitoring and Observability for Generative AI

- Foundation model monitoring systems
- Business impact and value management
- AI-specific troubleshooting and diagnostics

Module 9: Testing, Validation, and Continuous Improvement

- Comprehensive AI evaluation frameworks
- Quality assurance and continuous improvement
- RAG system evaluation and optimization

Module 10: Enterprise Integration Patterns

- Enterprise connectivity and integration architecture
- Secure access and identity management
- Cross-environment and hybrid deployments

Module 11: Course Wrap Up

6.2.8 Machine Learning Engineering on AWS

Course Description

Machine Learning (ML) Engineering on Amazon Web Services (AWS) is a 3-day intermediate course designed for ML professionals seeking to learn machine learning engineering on AWS. Participants learn to build, deploy, orchestrate, and operationalize ML solutions at scale through a balanced combination of theory, practical labs, and activities. Participants will gain practical experience using AWS services such as Amazon SageMaker AI and analytics tools such as Amazon EMR to develop robust, scalable, and production-ready machine learning applications.

- Course level: Intermediate
- Duration: 3 days

Course Objectives

In this course, you will learn to:

- Explain ML fundamentals and its applications in the AWS Cloud.
- Process, transform, and engineer data for ML tasks by using AWS services.
- Select appropriate ML algorithms and modeling approaches based on problem requirements and model interpretability.
- Design and implement scalable ML pipelines by using AWS services for model training, deployment, and orchestration.
- Create automated continuous integration and delivery (CI/CD) pipelines for ML workflows.
- Discuss appropriate security measures for ML resources on AWS.
- Implement monitoring strategies for deployed ML models, including techniques for detecting data drift.

Intended Audience

This course is designed for professionals who are interested in building, deploying, and operationalizing machine learning models on AWS. This could include current and in-training machine learning engineers who might have little prior experience with AWS. Other roles that can benefit from this training are DevOps engineer, developer, and SysOps engineer.

Prerequisites

We recommend that attendees of this course have the following:

- Familiarity with basic machine learning concepts
- Working knowledge of Python programming language and common data science libraries such as NumPy, Pandas, and Scikit-learn
- Basic understanding of cloud computing concepts and familiarity with AWS
- Experience with version control systems such as Git (beneficial but not required)

Machine Learning Engineering on AWS Course Outline

Day 1

Course Introduction

Module 1: Introduction to Machine Learning (ML) on AWS

- Introduction to ML
- Amazon SageMaker AI
- Responsible ML

Module 2: Analyzing Machine Learning (ML) Challenges

- Evaluating ML business challenges
- ML training approaches
- ML training algorithms

Module 3: Data Processing for Machine Learning (ML)

- Data preparation and types
- Exploratory data analysis
- AWS storage options and choosing storage

Module 4: Data Transformation and Feature Engineering

- Handling incorrect, duplicated, and missing data
- Feature engineering concepts
- Feature selection techniques
- AWS data transformation services
- Lab 1: Analyze and Prepare Data with Amazon SageMaker Data Wrangler and Amazon EMR
- Lab 2: Data Processing Using SageMaker Processing and the SageMaker Python SDK Day 2

Day 2

Module 5: Choosing a Modelling Approach

- Amazon SageMaker AI built-in algorithms
- Amazon SageMaker Autopilot
- Selecting built-in training algorithms
- Model selection considerations
- ML cost considerations



Module 6: Training Machine Learning (ML) Models

- Model training concepts
- Training models in Amazon SageMaker AI
- Lab 3: Training a model with Amazon SageMaker AI

Module 7: Evaluating and Tuning Machine Learning (ML) models

- Evaluating model performance
- Techniques to reduce training time
- Hyperparameter tuning techniques
- Lab 4: Model Tuning and Hyperparameter Optimization with Amazon SageMaker AI

Module 8: Model Deployment Strategies

- Deployment considerations and target options
- Deployment strategies
- Choosing a model inference strategy
- Container and instance types for inference
- Lab 5: Shifting Traffic

Day 3

Module 9: Securing AWS Machine Learning (ML) Resources

- Access control
- Network access controls for ML resources
- Security considerations for CI/CD pipelines

Module 10: Machine Learning Operations (MLOps) and Automated Deployment

- Introduction to MLOps
- Automating testing in CI/CD pipelines
- Continuous delivery services
- Lab 6: Using Amazon SageMaker Pipelines and the Amazon SageMaker Model Registry with Amazon SageMaker Studio

Module 11: Monitoring Model Performance and Data Quality

- Detecting drift in ML models
- SageMaker Model Monitor
- Monitoring for data quality and model quality
- Automated remediation and troubleshooting
- Lab 7: Monitoring a Model for Data Drift

Module 12: Course Wrap Up

6.2.9 MLOps Engineering on AWS

Course Description

This course builds upon and extends the DevOps methodology prevalent in software development to build, train, and deploy machine learning (ML) models. The course is based on the four-level MLOPs maturity framework. The course focuses on the first three levels, including the initial, repeatable, and reliable levels. The course stresses the importance of data, model, and code to successful ML deployments. It demonstrates 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 also discusses the use of tools and processes to monitor and act when the model prediction in production drifts from agreed-upon key performance indicators.

- Course level: Intermediate
- Duration: 3 days

Course Objectives

In this course, you will learn to:

- Explain the benefits of MLOps
- Compare and contrast DevOps and MLOps
- Evaluate the security and governance requirements for an ML use case and describe possible solutions and mitigation strategies
- Set up experimentation environments for MLOps with Amazon SageMaker
- Explain best practices for versioning and maintaining the integrity of ML model assets (data, model, and code)
- Describe three options for creating a full CI/CD pipeline in an ML context
- Recall best practices for implementing automated packaging, testing and deployment. (Data/model/code)
- Demonstrate how to monitor ML based solutions
- Demonstrate how to automate an ML solution that tests, packages, and deploys a model in an automated fashion; detects performance degradation; and re-trains the model on top of newly acquired data

Intended Audience

This course is intended for:

- MLOps engineers who want to produce and monitor ML models in the AWS cloud
- DevOps engineers who will be responsible for successfully deploying and maintaining ML models in production

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

MLOps Engineering on AWS Course Outline

Day 1

Module 1: Introduction to MLOps

- Processes
- People
- Technology
- Security and governance
- MLOps maturity model

Module 2: Initial MLOps: Experimentation Environments in SageMaker Studio

- Bringing MLOps to experimentation
- Setting up the ML experimentation environment
- Demonstration: Creating and Updating a Lifecycle Configuration for SageMaker Studio
- Hands-On Lab: Provisioning a SageMaker Studio Environment with the AWS Service Catalogue
- Workbook: Initial MLOps

Module 3: Repeatable MLOps: Repositories

- Managing data for MLOps
- Version control of ML models
- Code repositories in ML

Module 4: Repeatable MLOps: Orchestration

- ML pipelines
- Demonstration: Using SageMaker Pipelines to Orchestrate Model Building Pipelines

Day 2

Module 4: Repeatable MLOps: Orchestration (continued)

- End-to-end orchestration with AWS Step Functions
- Hands-On Lab: Automating a Workflow with Step Functions
- End-to-end orchestration with SageMaker Projects
- Demonstration: Standardizing an End-to-End ML Pipeline with SageMaker Projects
- Using third-party tools for repeatability
- Demonstration: Exploring Human-in-the-Loop During Inference
- Governance and security
- Demonstration: Exploring Security Best Practices for SageMaker
- Workbook: Repeatable MLOps

Module 5: Reliable MLOps: Scaling and Testing

- Scaling and multi-account strategies
- Testing and traffic-shifting
- Demonstration: Using SageMaker Inference Recommender



- Hands-On Lab: Testing Model Variants

Day 3

Module 5: Reliable MLOps: Scaling and Testing (continued)

- Hands-On Lab: Shifting Traffic
- Workbook: Multi-account strategies

Module 6: Reliable MLOps: Monitoring

- The importance of monitoring in ML
- Hands-On Lab: Monitoring a Model for Data Drift
- Operations considerations for model monitoring
- Remediating problems identified by monitoring ML solutions
- Workbook: Reliable MLOps
- Hands-On Lab: Building and Troubleshooting an ML Pipeline

6.2.10 Practical Data Science with Amazon SageMaker

Course Description

Artificial intelligence and machine learning (AI/ML) are becoming mainstream. In this course, you will spend a day in the life of a data scientist so that you can collaborate efficiently with data scientists and build applications that integrate with ML. You will learn the basic process data scientists use to develop ML solutions on Amazon Web Services (AWS) with Amazon SageMaker. You will experience the steps to build, train, and deploy an ML model through instructor-led demonstrations and labs.

- Course level: Intermediate
- Duration: 1 day

Course Objectives

In this course, you will learn how to:

- Discuss the benefits of different types of machine learning for solving business problems
- Describe the typical processes, roles, and responsibilities on a team that builds and deploys ML systems
- Explain how data scientists use AWS tools and ML to solve a common business problem
- Summarize the steps a data scientist takes to prepare data
- Summarize the steps a data scientist takes to train ML models
- Summarize the steps a data scientist takes to evaluate and tune ML models
- Summarize the steps to deploy a model to an endpoint and generate predictions
- Describe the challenges for operationalizing ML models
- Match AWS tools with their ML function

Intended Audience

This course is intended for DevOps Engineers and Application Developers.



Prerequisites

We recommend that attendees of this course have:

- AWS Technical Essentials
- Entry-level knowledge of Python programming
- Entry-level knowledge of statistics

Practical Data Science with Amazon SageMaker course outline

Module 1: Introduction to Machine Learning

- Benefits of machine learning (ML)
- Types of ML approaches
- Framing the business problem
- Prediction quality
- Processes, roles, and responsibilities for ML projects

Module 2: Preparing a Dataset

- Data analysis and preparation
- Data preparation tools
- Demonstration: Review Amazon SageMaker Studio and Notebooks
- Hands-On Lab: Data Preparation with SageMaker Data Wrangler

Module 3: Training a Model

- Steps to train a model
- Choose an algorithm
- Train the model in Amazon SageMaker
- Hands-On Lab: Training a Model with Amazon SageMaker
- Amazon Q Developer
- Demonstration: Amazon Q Developer in SageMaker Studio Notebooks

Module 4: Evaluating and Tuning a Model

- Model evaluation
- Model tuning and hyperparameter optimization
- Hands-On Lab: Model Tuning and Hyperparameter Optimization with Amazon SageMaker

Module 5: Deploying a Model

- Model deployment
- Hands-On Lab: Deploy a Model to a Real-Time Endpoint and Generate a Prediction

Module 6: Operational Challenges

- Responsible ML
- ML team and MLOps
- Automation
- Monitoring
- Updating models (model testing and deployment)



Module 7: Other Model-Building Tools

- Different tools for different skills and business needs
- No-code ML with Amazon SageMaker Canvas
- Demonstration: Overview of Amazon SageMaker Canvas
- Amazon SageMaker Studio Lab
- Demonstration: Overview of SageMaker Studio Lab
- (Optional) Hands-On Lab: Integrating a Web Application with an Amazon SageMaker Model Endpoint

6.2.11 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.

- Course level: Advanced
- Duration: 3 days

Course Objectives

In this course, you will learn to:

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

Intended Audience

This course is intended for experienced Data Scientists who are proficient in ML and deep learning fundamentals.

Prerequisites

We recommend that all attendees of this course have:

- Experience using ML frameworks
- Python programming experience
- At least 1 year of experience as a data scientist responsible for training, tuning, and deploying models
- AWS Technical Essentials training

Amazon SageMaker Studio for Data Scientists Course Outline

Day 1

Module 1: Amazon SageMaker Studio Setup

- JupyterLab Extensions in SageMaker Studio
- Demonstration: SageMaker user interface demo

Module 2: Data Processing

- Using SageMaker Data Wrangler for data processing
- Hands-On Lab: Analyse and prepare data using Amazon SageMaker Data Wrangler

- Using Amazon EMR
- Hands-On Lab: Analyse and prepare data at scale using Amazon EMR
- Using AWS Glue interactive sessions
- Using SageMaker Processing with custom scripts
- Hands-On Lab: Data processing using Amazon SageMaker Processing and SageMaker Python SDK
- SageMaker Feature Store
- Hands-On Lab: Feature engineering using SageMaker Feature Store

Module 3: Model Development

- SageMaker training jobs
- Built-in algorithms
- Bring your own script
- Bring your own container
- SageMaker Experiments
- Hands-On Lab: Using SageMaker Experiments to Track Iterations of Training and Tuning Models

Day 2

Module 3: Model Development (continued)

- SageMaker Debugger
- Hands-On Lab: Analysing, Detecting, and Setting Alerts Using SageMaker Debugger
- Automatic model tuning
- SageMaker Autopilot: Automated ML
- Demonstration: SageMaker Autopilot
- Bias detection
- Hands-On Lab: Using SageMaker Clarify for Bias and Explainability
- SageMaker Jumpstart

Module 4: Deployment and Inference

- SageMaker Model Registry
- SageMaker Pipelines
- Hands-On Lab: Using SageMaker Pipelines and SageMaker Model Registry with SageMaker Studio
- SageMaker model inference options
- Scaling
- Testing strategies, performance, and optimization
- Hands-On Lab: Inferencing with SageMaker Studio

Module 5: Monitoring

- Amazon SageMaker Model Monitor
- Discussion: Case study
- Demonstration: Model Monitoring



Day 3

Module 6: Managing SageMaker Studio Resources and Updates

- Accrued cost and shutting down
- Updates
- Environment setup
- Challenge 1: Analyse and prepare the dataset with SageMaker Data Wrangler
- Challenge 2: Create feature groups in SageMaker Feature Store
- Challenge 3: Perform and manage model training and tuning using SageMaker Experiments
- (Optional) Challenge 4: Use SageMaker Debugger for training performance and model optimization
- Challenge 5: Evaluate the model for bias using SageMaker Clarify
- Challenge 6: Perform batch predictions using model endpoint
- (Optional) Challenge 7: Automate full model development process using SageMaker Pipeline

6.3 Containers Course Curriculum

6.3.1 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.

- Course level: Intermediate
- Duration: 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.

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 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

6.4 Data Analytics Course Curriculum

6.4.1 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.

- Course level: Intermediate
- Duration: 1 day

Course Objectives

In this course, you will learn to:

- Apply data lake methodologies in planning and designing a data lake
- Plan and design a data lake using established data lake methodologies
- Describe the components and services required for building a data lake on AWS
- Secure a data lake with appropriate permissions
- Compare the ways data can be ingested, stored and transformed in a data lake on AWS
- Analyze and visualize data within a data lake
- Build and automate deployment of a data lake on AWS
- Describe the role of a data lake within modern data architecture

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, Cataloguing, 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 catalogue
- Identify data formatting, partitioning, and compression for efficient storage and query

Module 3: Building a Data Lake with AWS Lake Formation

- 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 analyse data in a data lake
- Lab 01: Building a Data Lake with AWS Lake Formation

Module 4: Data Processing and Analysis

- 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

- Explain the available built-in Blueprints to create and populate a new Lake Formation
- Describe methods for applying advanced permissions to secure data access and workflow.
- Describe fine-grained row/cell access control
- Explain the Lake Formation Tag-based access control mechanism and the different use cases for Named access control vs. Tag-based access control
- Describe access flow that enforces fine-grained access policies to both catalogue metadata and underlying data resource for analytics services connecting to Lake Formation

Module 6: Modern Data Architecture

- Explain capabilities of a modern data architecture: scalable data lakes, purpose-build analytics services, seamless data movement, unified governance, and performance and cost-effectiveness
- Articulate the typical data movement within a modern data architecture: Inside out, Outside in, Around the perimeter, and Sharing across.
- Describe focus of building and maintaining data products as a service.
- Describe a typical Data Mesh architecture using Lake Formation and the key enablers supporting this methodology
- Lab 3: Building and publishing a data product in Lake Formation

Module 7: Course Wrap Up

- Post course knowledge check
- Architecture review
- Course review

6.4.2 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.



- Course level: Intermediate
- Duration: 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
- 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.

Prerequisites

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
- Overview of Data Analytics and the Data Pipeline Data Analytics use cases
- Using the Data Pipeline for Data Analytics

Building Batch Data Analytics Solutions on AWS Course Outline

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 8: Developing Modern Data Architectures on AWS

- Modern Data Architectures

6.4.3 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.

- Course level: Intermediate
- Duration: 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

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 data 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 7: Developing Modern Data Architectures on AWS

- Modern Data Architectures

6.4.4 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.

- Course level: Intermediate
- Duration: 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, sharding 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

We recommend that attendees of this course have:

- 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 for those that need a refresher on streaming concepts.
- Completed either Architecting on AWS or Data Analytics Fundamentals
- Completed Building Data Lakes on AWS

Building Streaming Data Analytics Solutions on AWS Course Outline

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 8: Developing Modern Data Architectures on AWS

- Modern Data Architectures

6.4.5 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. This course demonstrates how to ingest, store, and transform data in the data warehouse. Topics covered include: the purpose of Amazon Redshift, how Amazon Redshift addresses business and technical challenges, features and capabilities of Amazon Redshift, designing a Data Warehousing Solution on AWS by applying best practices based on the Well-Architected Framework, integration with AWS and non-AWS products and services, performance tuning, orchestration, and securing and monitoring Amazon Redshift.



- Course level: Advanced
- Duration: 3 days

Course Objectives

In this course, you will learn to:

- Describe Amazon Redshift architecture and its roles in a modern data architecture
- Design and implement a data warehouse in the cloud using Amazon Redshift
- Identify and load data into an Amazon Redshift data warehouse from a variety of sources
- Analyse data using SQL QEV2 notebooks
- Design and implement a disaster recovery strategy for an Amazon Redshift data warehouse
- Perform maintenance and performance tuning on an Amazon Redshift data warehouse
- Secure and manage access to an Amazon Redshift data warehouse
- Share data between multiple Redshift clusters in an organisation
- Orchestrate workflows in the data warehouse using AWS Step Functions state machines
- Create an ML model and configure predictors using Amazon Redshift ML

Intended Audience

This course is intended for:

- Data Engineers
- Data Architects
- Database Architects
- Database Administrators
- Database Developers

Prerequisites

We recommend that attendees of this course have completed the following courses:

- Fundamentals of Analytics on AWS – Part 1 (Digital course)
- Fundamentals of Analytics on AWS – Part 2 (Digital course)
- Building Data Lakes on AWS (Instructor led Training)
- Building Data Analytics Solutions Using Amazon Redshift (Instructor led Training)

Data Warehousing on AWS Course Outline

Day 1

Module 1: Data Warehouse Concepts

- Modern data architecture
- Introduction to the course story
- Data warehousing with Amazon Redshift
- Amazon Redshift Serverless architecture
- Hands-On Lab: Launch and Configure an Amazon Redshift Serverless Data Warehouse



Module 2: Setting up Amazon Redshift

- Data models for Amazon Redshift
- Data management in Amazon Redshift
- Managing permissions in Amazon Redshift
- Hands-On Lab: Setting up a Data Warehouse using Amazon Redshift Serverless

Module 3: Loading Data

- Overview of data sources
- Loading data from Amazon Simple Storage Service (Amazon S3)
- Extract, transform, and load (ETL) and extract, load, and transform (ELT)
- Loading streaming data
- Loading data from relational databases
- Hands-On Lab: Populating the data warehouse

Day 2

Module 4: Deep Dive into SQL Query Editor v2 and Notebooks

- Features of Amazon Redshift Query Editor v2
- Demonstration: Using Amazon Redshift Query Editor v2
- Advanced queries
- Hands-On Lab: Data Wrangling on AWS

Module 5: Backup and Recovery

- Disaster recovery
- Backing up and restoring Amazon Redshift provisioned
- Backing up and restoring Amazon Redshift Serverless

Module 6: Amazon Redshift Performance Tuning

- Factors that impact query performance
- Table maintenance and materialized views
- Query analysis
- Workload management
- Tuning guidance
- Amazon Redshift monitoring
- Hands-On Lab: Performance Tuning the Data Warehouse

Module 7: Securing Amazon Redshift

- Introduction to Amazon Redshift security and compliance
- Authentication with Amazon Redshift
- Access control with Amazon Redshift
- Data encryption with Amazon Redshift
- Auditing and compliance with Amazon Redshift
- Hands-On Lab: Securing Amazon Redshift

Day 3



Module 8: Orchestration

- Overview of data orchestration
- Orchestration with AWS Step Functions
- Orchestration with Amazon Managed Workflows for Apache Airflow (MWAA)
- Hands-On Lab: Orchestrating the Data Warehouse Pipeline

Module 9: Amazon Redshift ML

- Machine Learning Overview
- Getting started with Amazon Redshift ML
- Amazon Redshift ML workflow scenarios
- Amazon Redshift ML Usage
- Hands-On Lab: Predicting customer churn with Amazon Redshift ML

Module 10: Amazon Redshift Data Sharing

- Overview of data sharing in Amazon Redshift
- Amazon DataZone for Data as a service

Module 11: Course Wrap Up

- Hands-On Lab: End of course challenge lab

6.4.6 Data Engineering on AWS

Course Description

Data Engineering on AWS is a 3-day intermediate course, designed for professionals seeking a deep dive into data engineering practices and solutions on AWS. Through a balanced combination of theory, practical labs, and activities, participants learn to design, build, optimize, and secure data engineering solutions using AWS services. From foundational concepts to hands-on implementation of data lakes, data warehouses, and both batch and streaming data pipelines, this course equips data professionals with the skills needed to architect and manage modern data solutions at scale.

- Course level: Intermediate
- Duration: 3 days

Course Objectives

In this course, you will learn to do the following:

- Understand the foundational roles and key concepts of data engineering, including data personas, data discovery, and relevant AWS services.
- Identify and explain the various AWS tools and services crucial for data engineering, encompassing orchestration, security, monitoring, CI/CD, IaC, networking, and cost optimization.
- Design and implement a data lake solution on AWS, including storage, data ingestion, transformation, and serving data for consumption.
- Optimize and secure a data lake solution by implementing open table formats, security measures, and troubleshooting common issues.
- Design and set up a data warehouse using Amazon Redshift Serverless, understanding its architecture, data ingestion, processing, and serving capabilities.

- Apply performance optimization techniques to data warehouses in Amazon Redshift, including monitoring, data optimization, query optimization, and orchestration.
- Manage security and access control for data warehouses in Amazon Redshift, understanding authentication, data security, auditing, and compliance.
- Design effective batch data pipelines using appropriate AWS services for processing and transforming data.
- Implement comprehensive strategies for batch data pipelines, covering data processing, transformation, integration, cataloguing, and serving data for consumption.
- Optimize, orchestrate, and secure batch data pipelines, demonstrating advanced skills in data processing automation and security.
- Architect streaming data pipelines, understanding various use cases, ingestion, storage, processing, and analysis using AWS services.
- Optimize and secure streaming data solutions, including compliance considerations and access control.

Intended Audience

This course is designed for professionals who are interested in designing, building, optimizing, and securing data engineering solutions using AWS services.

Prerequisites

We recommend that attendees of this course have:

- Familiarity with basic machine learning concepts, such as supervised and unsupervised learning, regression, classification, and clustering algorithms.
- Working knowledge of Python programming language and common data science libraries like NumPy, Pandas, and Scikit-learn.
- Basic understanding of cloud computing concepts and familiarity with the AWS platform.
- Familiarity with SQL and relational databases is recommended but not mandatory.

Data Engineering on AWS Course Outline

Day 1

Module 1: Data Engineering Roles and Key Concepts

- Role of a Data Engineer
- Key functions of a Data Engineer
- Data Personas
- Data Discovery
- AWS Data Services

Module 2: AWS Data Engineering Tools and Services

- Orchestration and Automation
- Data Engineering Security
- Monitoring
- Continuous Integration and Continuous Delivery
- Infrastructure as Code



- AWS Serverless Application Model
- Networking Considerations
- Cost Optimization Tools

Module 3: Designing and Implementing Data Lakes

- Data lake introduction
- Data lake storage
- Ingest data into a data lake
- Catalogue data
- Transform data
- Server data for consumption
- Hands-on lab: Setting up a Data Lake on AWS
- Open Table Formats
- Security using AWS Lake Formation
- Setting permissions with Lake Formation
- Security and governance
- Troubleshooting
- Hand-on lab: Automating Data Lake Creation using AWS Lake Formation Blueprints

Day 2

Module 5: Data Warehouse Architecture and Design Principles

- Introduction to data warehouses
- Amazon Redshift Overview
- Ingesting data into Redshift
- Processing data
- Serving data for consumption
- Hands-on Lab: Setting up a Data Warehouse using Amazon Redshift Serverless

Module 6: Performance Optimization Techniques for Data Warehouses

- Monitoring and optimization options
- Data optimization in Amazon Redshift
- Query optimization in Amazon Redshift
- Orchestration options

Module 7: Security and Access Control for Data Warehouses

- Authentication and access control in Amazon Redshift
- Data security in Amazon Redshift
- Auditing and compliance in Amazon Redshift
- Hands-on lab: Managing Access Control in Redshift

Module 8: Designing Batch Data Pipelines

- Introduction to batch data pipelines
- Designing a batch data pipeline
- AWS services for batch data processing



Module 9: Implementing Strategies for Batch Data Pipeline

- Elements of a batch data pipeline
- Processing and transforming data
- Integrating and cataloguing your data
- Serving data for consumption
- Hands-on lab: A Day in the Life of a Data Engineer

Day 3

Module 10: Optimizing, Orchestrating, and Securing Batch Data Pipelines

- Optimizing the batch data pipeline
- Orchestrating the batch data pipeline
- Securing the batch data pipeline
- Hands-on lab: Orchestrating Data Processing in Spark using AWS Step Functions

Module 11: Streaming Data Architecture Patterns

- Introduction to streaming data pipelines
- Ingesting data from stream sources
- Streaming data ingestion services
- Storing streaming data
- Processing Streaming Data
- Analysing Streaming Data with AWS Services
- Hands-on lab: Streaming Analytics with Amazon Managed Service for Apache Flink

Module 12: Optimizing and Securing Streaming Solutions

- Optimizing a streaming data solution
- Securing a streaming data pipeline
- Compliance considerations
- Hands-on lab: Access Control with Amazon Managed Streaming for Apache Kafka

6.5 Databases Course Curriculum

6.5.1 Build Modern 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.

- Course level: Intermediate
- Duration: 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 Q Developer

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

6.6 Developer Course Curriculum

6.6.1 Developing on AWS

Course Description

This course teaches experienced developers how to programmatically interact with AWS services to build web solutions. It guides you through a high-level architectural discussion on resource selection and dives deep into using the AWS Software Development Kits (AWS SDKs) and Command Line Interface (AWS CLI) to build and deploy your cloud applications. You will build a sample application during this course, learning how to set up permissions to the development environment, adding business logic to process data using AWS core services, configure user authentications, deploy to AWS cloud, and debug to resolve application issues. The course includes code examples to help you implement the design patterns and solutions discussed in the course. The labs reinforce key course content and

help you to implement solutions using the AWS SDK for Python, .Net, and Java, the AWS CLI, and the AWS Management Console.

- Course level: Intermediate
- Duration: 3 days

Course Objectives

In this course, you will learn to:

- Build a simple end-to-end cloud application using AWS Software Development Kits (AWS SDKs), Command Line Interface (AWS CLI), and IDEs.
- Configure AWS Identity and Access Management (IAM) permissions to support a development environment.
- Use multiple programming patterns in your applications to access AWS services.
- Use AWS SDKs to perform CRUD (create, read, update, delete) operations on Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB resources.
- Build AWS Lambda functions with other service integrations for your web applications.
- Understand the benefits of microservices architectures and serverless applications to design.
- Develop API Gateway components and integrate with other AWS services.
- Explain how Amazon Cognito controls user access to AWS resources.
- Build a web application using Cognito to provide and control user access.
- Use DevOps methodology to reduce the risks associated with traditional application releases and identify AWS services that help in implementing DevOps practices.
- Use AWS Serverless Application Model (AWS SAM) to deploy an application.
- Observe your application build using Amazon X-Ray.

Intended Audience

This course is intended for:

- Software developers
- Solutions Architects
- IT workers who want to improve their developing skills using AWS Services

Prerequisites

We recommend that attendees of this course have:

- AWS Technical Essentials
- Working knowledge of AWS core services
- Programming experience in any one of the following languages:
 - Python
 - .NET
 - Java

Developing on AWS Course Outline

Day 1

Module 1: Course Overview

- Logistics
- Student resources
- Agenda
- Introductions

Module 2: Building a Web Application on AWS

- Discuss the architecture of the application you are going to build during this course
- Explore the AWS services needed to build your web application
- Discover how to store, manage, and host your web application

Module 3: Getting Started with Development on AWS

- Describe how to access AWS services programmatically
- List some programmatic patterns and how they provide efficiencies within AWS SDKs and AWS CLI
- Explain the value of AWS Cloud9

Module 4: Getting Started with Permissions

- Review AWS Identity and Access Management (IAM) features and components permissions to support a development environment
- Demonstrate how to test AWS IAM permissions
- Configure your IDEs and SDKs to support a development environment
- Demonstrate accessing AWS services using SDKs and AWS Cloud9
- Lab 1: Configure the Developer Environment
- Connect to a developer environment
- Verify that the IDE and the AWS CLI are installed and configured to use the application profile
- Verify that the necessary permissions have been granted to run AWS CLI commands
- Assign an AWS IAM policy to a role to delete an Amazon S3 bucket

Module 5: Getting Started with Storage

- Describe the basic concepts of Amazon S3
- List the options for securing data using Amazon S3
- Define SDK dependencies for your code
- Explain how to connect to the Amazon S3 service
- Describe request and response objects

Module 6: Processing Your Storage Operations

- Perform key bucket and object operations
- Explain how to handle multiple and large objects
- Create and configure an Amazon S3 bucket to host a static website
- Grant temporary access to your objects



- Demonstrate performing Amazon S3 operations using SDKs
- Lab 2: Develop Solutions Using Amazon S3
- Interact with Amazon S3 programmatically using AWS SDKs and the AWS CLI
- Create a bucket using waiters and verify service exceptions codes
- Build the needed requests to upload an Amazon S3 object with metadata attached
- Build requests to download an object from the bucket, process data, and upload the object back to the bucket
- Configure a bucket to host the website and sync the source files using the AWS CLI
- Add IAM bucket policies to access the S3 website.

Day 2

Module 7: Getting Started with Databases

- Describe the key components of DynamoDB
- Explain how to connect to DynamoDB
- Describe how to build a request object
- Explain how to read a response object
- List the most common troubleshooting exceptions

Module 8: Processing Your Database Operations

- Develop programs to interact with DynamoDB using AWS SDKs
- Perform CRUD operations to access tables, indexes, and data
- Describe developer best practices when accessing DynamoDB
- Review caching options for DynamoDB to improve performance
- Perform DynamoDB operations using SDK
- Lab 3: Develop Solutions Using Amazon DynamoDB
- Interact with Amazon DynamoDB programmatically using low-level, document, and high-level APIs in your programs
- Retrieve items from a table using key attributes, filters, expressions, and paginations
- Load a table by reading JSON objects from a file
- Search items from a table based on key attributes, filters, expressions, and paginations
- Update items by adding new attributes and changing data conditionally
- Access DynamoDB data using PartiQL and object-persistence models where applicable

Module 9: Processing Your Application Logic

- Develop a Lambda function using SDKs
- Configure triggers and permissions for Lambda functions
- Test, deploy, and monitor Lambda functions
- Lab 4: Develop Solutions Using AWS Lambda Functions
- Create AWS Lambda functions and interact programmatically using AWS SDKs and AWS CLI
- Configure AWS Lambda functions to use the environment variables and to integrate with other services



- Generate Amazon S3 pre-signed URLs using AWS SDKs and verify the access to bucket objects
- Deploy the AWS Lambda functions with .zip file archives through your IDE and test as needed
- Invoke AWS Lambda functions using the AWS Console and AWS CLI

Module 10: Managing the APIs

- Describe the key components of API Gateway
- Develop API Gateway resources to integrate with AWS services
- Configure API request and response calls for your application endpoints
- Test API resources and deploy your application API endpoint
- Demonstrate creating API Gateway resources to interact with your application APIs
- Lab 5: Develop Solutions Using Amazon API Gateway
- Create RESTful API Gateway resources and configure CORS for your application
- Integrate API methods with AWS Lambda functions to process application data
- Configure mapping templates to transform the pass-through data during method integration
- Create a request model for API methods to ensure that the pass-through data format complies with application rules
- Deploy the API Gateway to a stage and validate the results using the API endpoint

Day 3

Module 11: Building a Modern Application

- Describe the challenges with traditional architectures
- Describe the microservice architecture and benefits
- Explain various approaches for designing microservice applications
- Explain steps involved in decoupling monolithic applications
- Demonstrate the orchestration of Lambda Functions using AWS Step Functions

Module 12: Granting Access to Your Application Users

- Analyse the evolution of security protocols
- Explore the authentication process using Amazon Cognito
- Manage user access and authorize serverless APIs
- Observe best practices for implementing Amazon Cognito
- Demonstrate the integration of Amazon Cognito and review JWT tokens

Lab 6: Capstone – Complete the Application Build

- Create a Userpool and an Application Client for your web application using
- Add new users and confirm their ability to sign-in using the Amazon Cognito CLI
- Configure API Gateway methods to use Amazon Cognito as an authorizer
- Verify JWT authentication tokens are generated during API Gateway calls
- Develop API Gateway resources rapidly using a Swagger importing strategy
- Set up your web application frontend to use Amazon Cognito and API Gateway configurations and verify the entire application functionality

Module 13: Deploying Your Application

- Identify risks associated with traditional software development practices
- Understand DevOps methodology
- Configure an AWS SAM template to deploy a serverless application
- Describe various application deployment strategies
- Demonstrate deploying a serverless application using AWS SAM

Module 14: Observing Your Application

- Differentiate between monitoring and observability
- Evaluate why observability is necessary in modern development and key components
- Understand CloudWatch's part in configuring the observability
- Demonstrate using CloudWatch Application Insights to monitor applications
- Demonstrate using X-Ray to debug your applications
- Lab 7: Observe the Application Using AWS X-Ray
- Instrument your application code to use AWS X-Ray capabilities
- Enable your application deployment package to generate logs
- Understand the key components of an AWS SAM template and deploy your application
- Create AWS X-Ray service maps to observe end-to-end processing behaviour of your application
- Analyse and debug application issues using AWS X-Ray traces and annotations

Module 15: Course Wrap Up

6.7 DevOps Course Curriculum

6.7.1 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 organisation'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.

- Course level: Intermediate
- Duration: 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 programming 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 1

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 2

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 3

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

Module 14: Course Summary

6.8 Migrate and Transfer Course Curriculum

6.8.1 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.

- Course level: Fundamental
- Duration: 1 day

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

AWS Migration Essentials Course Outline

Course Introduction

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

6.8.2 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 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 organisation.

- Course level: Intermediate
- Duration: 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 organisations 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 1

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 Organisations
- 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 policies
- 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 2

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 3

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

6.9 Networking and Content Delivery Course Curriculum

6.9.1 Networking Essentials for Cloud Applications on AWS

Course Description

This course provides a comprehensive understanding of networking concepts and services within Amazon Web Services Cloud environment. Designed for novice and experienced 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.

- Course level: Intermediate
- Duration: 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



6.10 Operations Course Curriculum

6.10.1 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

- Course level: Intermediate
- Duration: 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 Organisations 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 an 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 1

Module 1: Introduction to Cloud Operations on AWS

- What are Cloud Operations?
- AWS Well-Architected Framework
- AWS Well-Architected Tool

Module 2: Access Management

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

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 2

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 manage 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

Module 15: Course Wrap Up

6.11 Serverless Course Curriculum

6.11.1 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.

- Course level: Intermediate
- Duration: 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 organisation 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 1

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 2

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

Day 3

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

Module 13: Course Wrap Up

6.12 Security Course Curriculum

6.12.1 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 organisation.

- Course level: Fundamental
- Duration: 1 day

Course Objectives

In this course, you will learn to:

- Identify security benefits and responsibilities of using the AWS Cloud.
- Describe the access control and management features of AWS.
- Explain the available methods for encrypting data at rest and in transit.
- Describe how to secure network access to your AWS resources.
- Determine which AWS services can be used for monitoring and incident response

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

Course Introduction

Module 1: Exploring the Security Pillar

- AWS Well-Architected Framework: Security Pillar

Module 2: Security of the Cloud

- Shared responsibility model
- AWS Global Infrastructure
- Compliance and governance

Module 3: Identity and Access Management

- Identity and access management
- Data access and protection essentials
- Lab 1: Introduction to Security Policies

Module 4: Protecting Infrastructure and Data

- Protecting your network infrastructure
- Edge Security
- DDoS Mitigation
- Protecting compute resources
- Lab 2: Securing VPC Resources with Security Groups

Module 5: Detection and Response

- Monitoring and detective controls
- Incident response essentials

Module 6: Course Wrap Up

6.12.2 Security Engineering on AWS

Course Description

Security is a concern for both customers in the cloud, and those considering cloud adoption. An increase in cyberattacks and data leaks remains top of mind for most industry personnel. The Security Engineering on AWS course addresses these concerns by helping you better understand how to interact and build with Amazon Web Services (AWS) in a secure way. In this course, you will learn about managing identities and roles, managing and provisioning accounts, and monitoring API activity for anomalies. You will also learn about how to protect data stored on AWS. The course explores how you can generate, collect, and monitor logs

to help identify security incidents. Finally, you will review detecting and investigating security incidents with AWS services.

- Course level: Intermediate
- Duration: 3 days

Course Objectives

In this course you will:

- State an understanding of AWS cloud security based on the CIA triad.
- Create and analyze authentication and authorizations with IAM.
- Manage and provision accounts on AWS with appropriate AWS services.
- Identify how to manage secrets using AWS services.
- Monitor sensitive information and protect data via encryption and access controls.
- Identify AWS services that address attacks from external sources.
- Monitor, generate, and collect logs. • Identify indicators of security incidents.
- Identify how to investigate threats and mitigate using AWS services. 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
- Cloud architects
- Cloud operators working across all global segments

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 1

Module 1: Security Overview and Review

- Explain Security in the AWS Cloud.
- Explain AWS Shared Responsibility Model.
- Summarize IAM, Data Protection, and Threat Detection and Response.
- State the different ways to interact with AWS using the console, CLI, and SDKs.
- Describe how to use MFA for extra protection.
- State how to protect the root user account and access keys.

Module 2: Securing Entry Points on AWS

- Describe how to use multi-factor authentication (MFA) for extra protection.
- Describe how to protect the root user account and access keys.
- Describe IAM policies, roles, policy components, and permission boundaries.
- Explain how API requests can be logged and viewed using AWS CloudTrail and how to view and analyze access history.
- Hands-On Lab: Using Identity and Resource Based Policies

Module 3: Account Management and Provisioning on AWS

- Explain how to manage multiple AWS accounts using AWS Organisations and AWS Control Tower
- Explain how to implement multi-account environments with AWS Control Tower
- Demonstrate the ability to use identity providers and brokers to acquire access to AWS services
- Explain the use of AWS IAM Identity Center (successor to AWS Single Sign-On) and AWS Directory Service. Demonstrate the ability to manage domain user access with Directory Service and IAM Identity Center
- Hands-On Lab: Managing Domain User Access with AWS Directory Service

Day 2

Module 4: Secrets Management on AWS

- Describe and list the features of AWS KMS, CloudHSM, AWS Certificate Manager (ACM), and AWS Secrets Manager
- Demonstrate how to create a multi-Region AWS KMS key
- Demonstrate how to encrypt a Secrets Manager secret with an AWS KMS key
- Demonstrate how to use an encrypted secret to connect to an Amazon Relational Database Service (Amazon RDS) database in multiple AWS Regions
- Hands-on lab: Using AWS KMS to Encrypt Secrets in Secrets Manager

Module 5: Data Security

- Monitor data for sensitive information with Amazon Macie.
- Describe how to protect data at rest through encryption and access controls.
- Identify AWS services used to replicate data for protection.
- Determine how to protect data after it has been archived.
- Hands-on lab: Data Security in Amazon S3

Module 6: Infrastructure Edge Protection

- Describe the AWS features used to build secure infrastructure
- Describe the AWS services used to create resiliency during an attack
- Identify the AWS services used to protect workloads from external threats
- Compare the features of AWS Shield and AWS Shield Advanced
- Explain how centralized deployment for AWS Firewall Manager can enhance security
- Hands-on lab: Using AWS WAF to Mitigate Malicious Traffic

Day 3

Module 7: Monitoring and Collecting Logs on AWS

- Identify the value of generating and collecting logs
- Use Amazon Virtual Private Cloud (Amazon VPC) Flow Logs to monitor for security events
- Explain how to monitor for baseline deviations. • Describe Amazon EventBridge events
- Describe Amazon CloudWatch metrics and alarms
- List log analysis options and available techniques
- Identify use cases for using virtual private cloud (VPC) Traffic Mirroring
- Hands-on lab: Monitoring for and Responding to Security Incidents

Module 8: Responding to Threats

- Classify incident types in incident response
- Understand incident response workflows
- Discover sources of information for incident response using AWS services
- Understand how to prepare for incidents
- Detect threats using AWS services
- Analyze and respond to security findings
- Hands-on lab: Incident Response

6.13 Storage Course Curriculum

6.13.1 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.

- Course level: Intermediate
- Duration: 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 1

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 2

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 OpenZFS
- Amazon FSx for Lustre
- Choosing an Amazon FSx service
- Hands-On Lab: Working with FSx for NetApp ONTAP and FSx for OpenZFS

Day 3

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

7 Exam Readiness

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