

INT364: CLOUD ARCHITECTURE AND IMPLEMENTATION-II

L:2 T:0 P:2 Credits:3

Course Outcomes: Through this course students should be able to

CO1 :: describe AWS-native disaster recovery services and recovery architecture patterns for ensuring data durability and operational resilience.

CO2 :: cognize secure and scalable networks using VPC, Transit Gateway, VPN, and related connectivity services

CO3 :: explore storage and compute solutions with Amazon EC2, S3, and associated services for optimized performance

CO4 :: practice relational and NoSQL databases on AWS, including deployment, backup, and migration strategies

CO5 :: assess availability and responsiveness through monitoring, load balancing, auto scaling, and infrastructure automation

CO6 :: demonstrate architectural considerations to integrate serverless and container-based solutions using AWS Lambda, API Gateway and ECS/EKS

Unit I

Cloud Fundamentals and Security : AWS-native disaster recovery mechanisms and service orchestration for data durability and operational resilience, recovery architecture patterns - Backup & Restore, Pilot Light, Warm Standby, Multi-Site(Active-Active) configurations

Unit II

Networking and Connectivity in AWS : Amazon VPC and subnet design, Securing network resources using security groups and NACLs, Connecting to managed AWS services within a VPC, VPC peering and AWS Transit Gateway, Remote connectivity with Site-to-Site VPN and AWS Direct Connect, Monitoring networks and applying Well-Architected principles

Unit III

Storage and Compute Services : Amazon S3 fundamentals and storage classes, Data transfer, lifecycle policies, versioning in S3, Launching EC2 instances and choosing AMIs, EC2 instance types, user data, and configuration, Storage options like EBS and EFS for EC2, Applying Well-Architected Framework to storage and compute layers

Unit IV

Databases and Data Management : Database layer considerations and use cases, Amazon RDS features and connection management, Automated backups and read replicas in RDS, Amazon DynamoDB and other purpose-built databases, Migrating databases to AWS, Applying Well-Architected principles to the database layer

Unit V

Resiliency, Monitoring, and Automation : Monitoring AWS resources with CloudWatch, Scaling compute resources using Auto Scaling, Using load balancers for high availability, DNS-based routing with Amazon Route 53, Infrastructure automation with AWS CloudFormation, Automating deployments with AWS Quick Starts and Amazon Q Developer

Unit VI

Serverless and Microservices Architectures : serverless and architectural considerations, Building serverless applications with AWS Lambda,, Building serverless applications with AWS Lambda, Extending functionality using Amazon API Gateway, Running microservices with AWS container services, Applying Well-Architected principles to serverless architectures

List of Practicals / Experiments:

List of Practicals

- Exploring AWS Identity and Access Management (IAM)
- Hands on practice of Amazon S3 Transfer Acceleration, Managing lifecycles in Amazon S3, Amazon S3 versioning
- Creating a Static Website for the Cafe
- Hands on practice for Amazon Elastic File System (Amazon EFS)

- Configuring an EC2 Instance with User Data
- Creating a Dynamic Website for the Café
- Creating an Amazon RDS Database, Amazon RDS Automated Backup and Read Replicas
- Creating a VPC Networking Environment for the Café
- Creating a VPC Peering Connection
- Securing Applications by using Amazon Cognito
- Creating a Scalable and Highly Available Environment for the Café
- Implementing a Serverless Architecture on AWS

References: 1. WELCOME TO AWS DOCUMENTATION by AMAZON, NA