



Free Questions for Cloud-Deployment-and-Operations

Shared by Leon on 16-04-2026

For More Free Questions and Preparation Resources

[Check the Links on Last Page](#)



Question 1

Question Type: MultipleChoice

(Which service should be used to schedule the patching of EC2 instances and on-premises servers with security updates?)

Options:

- A- EventBridge
- B- Config
- C- App Runner
- D- Systems Manager



Answer:

D

Explanation:

AWS Systems Manager should be used to schedule the patching of EC2 instances and on-premises servers with security updates, utilizing features like Patch Manager and maintenance windows. This service supports hybrid environments and automated patch deployment. The WGU Cloud Deployment and Operations Study Guide (Section 5.2, Systems Manager Patch Manager) states, 'Systems Manager provides Patch Manager to schedule and deploy security updates to EC2 instances and on-premises servers during maintenance windows, ensuring consistent patching across hybrid environments.' EventBridge, Config, and App Runner do not offer this patching capability.



Question 2

Question Type: MultipleChoice

(Which two solutions should an administrator use to receive emails when a Lambda function returns an error? Select 2 answers.)

Options:

- A- Amazon Simple Queue Service

- B- Amazon Simple Notification Service
- C- Amazon CloudWatch
- D- AWS CloudTrail

Answer:

B, C

Explanation:

To receive emails when a Lambda function returns an error, the administrator should use Amazon CloudWatch to monitor the function's logs and metrics (e.g., errors) and Amazon Simple Notification Service (SNS) to send email notifications based on CloudWatch alarms. The WGU Cloud Deployment and Operations Study Guide (Section 4.3, CloudWatch and SNS) states, 'CloudWatch can detect Lambda errors via logs and trigger an alarm, which integrates with SNS to send email notifications to subscribed endpoints, ensuring timely error alerts.' SQS and CloudTrail are not designed for this notification workflow.

Question 3

Question Type: MultipleChoice

(A company builds an application that renders 3D movies for its users. The application does this by spawning a fleet of instances that each take a slice of the workload and then render a few hundred frames in around 30 minutes. If processing on a node is interrupted, the work can be moved to other running nodes. The cost to users is based on the underlying AWS cost plus a margin. Which EC2 instance type is most suitable to run workload pattern at the lowest cost?)

Options:

- A- Spot Instance
- B- On-Demand Instance
- C- Dedicated Instance
- D- Reserved Instance

Answer:

A

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

Spot Instances are the most suitable EC2 instance type for this workload pattern, offering the lowest cost for short, interruptible tasks like 3D rendering, where work can be redistributed if interrupted. This aligns with the company's cost-plus pricing model. The WGU Cloud Deployment and Operations Study Guide (Section 7.3, EC2 Instance Types) states, 'Spot Instances provide significant cost savings (up to 90% off On-Demand) for fault-tolerant, short-term workloads like rendering, where interruptions can be handled by redistributing tasks to other nodes.' On-Demand, Dedicated, and Reserved Instances are more expensive and less flexible for this use case.



Question 4

Question Type: MultipleChoice

(What is the advantage of maximizing the time to live of the cache in CloudFront?)

Options:

- A- Minimize requests to the origin
- B- Facilitate a high number of requests
- C- Improve a disaster recovery plan
- D- Improve security measures

Answer:

A



Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

Maximizing the time to live (TTL) of the cache in Amazon CloudFront reduces the frequency of requests to the origin server by serving content from edge locations for a longer period, lowering latency and costs. The WGU Cloud Deployment and Operations Study Guide (Section 4.4, CloudFront Caching) states, 'A higher TTL in CloudFront caching minimizes origin requests by keeping objects in edge caches longer, reducing load on the origin server and optimizing performance and cost.' Facilitating high request volumes, improving disaster recovery, and enhancing security are not primary benefits of maximizing TTL.

Question 5

Question Type: MultipleChoice

(An administrator needs to set up self-managed permissions for AWS CloudFormation StackSet operations. Which two roles should be created in the administrator account and in the target accounts? Choose 2 answers.)

Options:

- A- In each target account, create a service role named `AWSCloudFormationStackSetExecutionRole` that trusts the administrator account.
- B- In each target account, create a service role named `AWSCloudFormationStackSetAdministrationRole` that trusts the administrator account.
- C- In the administrator account, create an IAM role named `AWSCloudFormationStackSetAdministrationRole`.
- D- In the administrator account, create an IAM role named `AWSCloudFormationStackSetExecutionRole`.

Answer:

A, C

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

To set up self-managed permissions for CloudFormation StackSet operations, two roles are required: in the administrator account, create an IAM role named `AWSCloudFormationStackSetAdministrationRole` to manage StackSet operations, and in each target account, create a service role named `AWSCloudFormationStackSetExecutionRole` that trusts the administrator account to execute the stacks. The WGU Cloud Deployment and Operations Study Guide (Section 5.4, StackSets) states, 'For self-managed permissions, the administrator account requires the `AWSCloudFormationStackSetAdministrationRole`, while each target account needs the `AWSCloudFormationStackSetExecutionRole` with a trust relationship to the administrator account for stack deployment.' Roles B and D are incorrectly named or misplaced.

Question 6

Question Type: MultipleChoice

(Media files are stored in an S3 bucket in region A. Demand for the files from region B has increased. Region B is suffering from some latency issues due to geographical location. Which type of replication will reduce the latency in region B?)

Options:

- A- S3 Batch
- B- S3 Versioning
- C- Cross-Region
- D- Same-Region

Answer:

C

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

To reduce latency in region B due to increased demand for media files stored in region A, cross-region replication (CRR) should be used. CRR replicates objects to an S3 bucket in another region (e.g., region B), allowing users to access files from a closer location, thus minimizing latency. The WGU Cloud Deployment and Operations Study Guide (Section 2.2, S3 Replication) states, 'Cross-Region Replication in S3 copies objects to a destination bucket in a different region, reducing latency by enabling access to data from the nearest region, such as region B in this case.' S3 Batch, S3 Versioning, and Same-Region replication do not address cross-region latency.

Question 7

Question Type: MultipleChoice

(What should be used to monitor estimated AWS charges?)

Options:

- A- Cognito
- B- Billing alarm
- C- Pricing Calculator
- D- Forecast

Answer:

B

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

A Billing alarm in Amazon CloudWatch should be used to monitor estimated AWS charges by setting thresholds on cost metrics and triggering notifications when exceeded. This helps manage budgets effectively. The WGU Cloud Deployment and Operations Study Guide (Section 7.4, Cost Monitoring) states, 'Billing alarms in CloudWatch monitor estimated charges against a defined threshold, sending alerts via SNS to prevent cost overruns.' Cognito handles user authentication, Pricing Calculator estimates costs pre-deployment, and Forecast predicts usage, none of which monitor real-time charges.

Question 8

Question Type: MultipleChoice

(Which solution should be used to host content to be processed for Amazon Made?)

Options:

- A- DocumentDB
- B- DynamoDB
- C- S3
- D- EC2

Answer:

C

Explanation:

Amazon S3 (Simple Storage Service) should be used to host content to be processed for Amazon Made, as it provides scalable object storage ideal for storing and retrieving large amounts of data, such as media files or documents, for processing workflows. The WGU Cloud Deployment and Operations Study Guide (Section 2.1, Amazon S3) states, 'S3 is the preferred solution for hosting content to be processed by services like Amazon Made, offering durable and highly available storage with support for lifecycle policies and integration with other AWS services.' DocumentDB, DynamoDB, and EC2 are not designed for this content hosting purpose.



To Get Premium Files for Cloud-Deployment-and-Operations Visit

<https://www.p2pexams.com/products/cloud-deployment-and-operations>

For More Free Questions Visit

<https://www.p2pexams.com/wgu/pdf/cloud-deployment-and-operations>

20%
DISCOUNT

P2P
exams