WGU Cloud-Deployment-and-Operations D341 Practice Questions

Shared by Watts on 30-10-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

(What is used to change stacks across multiple accounts and Regions in a single operation?)

Options:

A- StackSets
B- StackInstance
C- Nested stacks
D- Stack policies

P2P exams

Answer:

A

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

StackSets are used to change stacks across multiple AWS accounts and regions in a single operation, enabling centralized management of infrastructure deployments. StackSets allow administrators to create, update, or delete stacks consistently across specified accounts and regions. The WGU Cloud Deployment and Operations Study Guide (Section 5.4, StackSets) states, 'StackSets provide the capability to manage and update stacks across multiple accounts and regions with a single operation, streamlining multi-region deployments.' StackInstance, nested stacks, and stack policies do not support this multi-account, multi-region functionality.

P2P exams

Question 2

Question Type: MultipleChoice

(What is the role of a patch baseline in Patch Manager?)

Options:

A- Installs all patch updates as they become available by default
B- Assigns an IAM role for services that can auto-approve patches upon release
C- Defines patches that should and should not be installed on EC2 instances

D- Notifies Config of any patch updates that need to be performed

Answer:

C

Explanation:

A patch baseline in Patch Manager defines the patches that should and should not be installed on EC2 instances, specifying approval rules for security updates, other updates, and exclusions. This ensures controlled and compliant patching. The WGU Cloud Deployment and Operations Study Guide (Section 5.2, Patch Manager) states, 'A patch baseline specifies which patches are approved for installation and which are excluded, allowing customization of patch deployment (e.g., approving only critical security updates).' Options A, B, and D misrepresent the baseline's role, as it does not auto-install, assign IAM roles, or notify Config directly.

Question 3

Question Type: MultipleChoice

(Which AWS solution can be used to send events from Shopify?)

Options:

A- EventBridge
B- CloudTrail events
C- IoT Events
D- Service events

Answer:

A

Explanation:

Amazon EventBridge is the AWS solution that can be used to send events from external sources like Shopify. EventBridge supports integration with third-party applications through its event bus, allowing custom events (e.g., from Shopify via webhooks) to be ingested and routed to AWS services or targets. The WGU Cloud Deployment and Operations Study Guide (Section 4.3, EventBridge) notes that EventBridge is designed for event ingestion from SaaS applications, making it suitable for Shopify integration. CloudTrail, IoT Events, and Service events are not

designed for this purpose.

Question 4

Question Type: MultipleChoice

(An AWS SysOps administrator needs to configure Amazon Route 53 to load balance customer requests across multiple identically configured websites. The load balancing method must ensure that the customer requests are directed to the website that offers the shortest round-trip time. Which routing policy should be used?)

Options:

A- Latency
B- Geolocation
C- Geoproximity
D- Weighted

Answer:

A

Explanation:

The latency routing policy in Amazon Route 53 should be used to direct customer requests to the website with the shortest round-trip time (RTT). This policy routes traffic to the AWS region with the lowest latency based on the user's location, optimizing performance. The WGU Cloud Deployment and Operations Study Guide (Section 3.1, Route 53 Routing Policies) states, 'The latency-based routing policy in Route 53 measures RTT to direct users to the region with the lowest latency, improving user experience across multiple websites.' Geolocation, geoproximity, and weighted policies do not prioritize RTT in this manner.

Question 5

Question Type: MultipleChoice

(A company stores critical data on general-purpose Amazon EBS volumes. The company can only tolerate the loss of up to one hour of data. Which solution should be used to protect data

from loss?)

Options:

A- Replace existing volumes with local storage
B- Schedule automated volume snapshots using CloudWatch Events
C- Enable automated volume backups in Cloud Control
D- Switch to Block Express volume type

Answer:

B

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

To protect critical data on EBS volumes with a maximum tolerable data loss of one hour, the company should schedule automated volume snapshots using CloudWatch Events. Snapshots capture the state of the volume at a point in time, and regular scheduling (e.g., every hour) ensures data loss is limited to the snapshot interval. The WGU Cloud Deployment and Operations Study Guide (Section 7.3, EBS and Snapshots) states, 'Automated EBS snapshots can be scheduled using CloudWatch Events rules to create backups at defined intervals, ensuring an RPO of one hour by capturing volume data regularly.' Local storage, Cloud Control backups, and Block Express are not viable solutions for this automated snapshot requirement.

Question 6

Question Type: MultipleChoice

(Which component should be used to configure email notifications for a trigger by Amazon Made?)

Options:

A- Simple Queue Service
B- Elastic Kubernetes Service
C- EventBridge
D- CodeDeploy

Answer:

C

Explanation:

Amazon EventBridge should be used to configure email notifications for a trigger by Amazon Made, as it can route events to targets like Amazon SNS (Simple Notification Service) for email delivery. EventBridge integrates with AWS services to handle event-driven notifications. The WGU Cloud Deployment and Operations Study Guide (Section 4.3, EventBridge) states, 'EventBridge can be configured to send events from Amazon Made to SNS, enabling email notifications by subscribing an email endpoint to the SNS topic.' SQS, EKS, and CodeDeploy do not support this notification configuration.

Question 7

Question Type: MultipleChoice

(A company configures the CloudFront caching TTL to the following settings: Min: 0, Max: 172800. An administrator updates the Cache-Control setting to 432000. However, objects are only cached for 48 hours. What is required to ensure objects are cached for up to 120 hours?)

Options:

A- Configure the Max setting for 432000
B- Configure the Min setting for 432000
C- Configure the expiration setting for 172800
D- Configure the Cache-Control setting for 172800

Answer:

A

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

In Amazon CloudFront, the caching TTL is controlled by the minimum and maximum TTL settings, with the maximum TTL (in seconds) capping the cache duration. The current Max TTL is 172800 seconds (48 hours), and the Cache-Control setting of 432000 seconds (120 hours) is ignored because it exceeds the Max TTL. To cache objects for up to 120 hours, the Max setting

must be configured to 432000 seconds. The WGU Cloud Deployment and Operations Study Guide (Section 4.4, CloudFront Caching) states, 'The maximum TTL in CloudFront settings (e.g., 172800 seconds) overrides any higher Cache-Control value; to allow 120 hours (432000 seconds), the Max TTL must be set accordingly.' Adjusting Min, expiration, or Cache-Control alone does not resolve this.

Question 8

Question Type: MultipleChoice

(An organization uses CloudFormation to deploy AWS infrastructure. The templates are built in JSON and deploy EC2 instances across multiple regions. Which solution should be used to set values in the template based on region-specific AMI IDs?)

Options:

A- Outputs
B- WaitCondition
C- Mappings
D- Resources

Answer:

C

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

The Mappings section in a CloudFormation template should be used to set values based on region-specific AMI IDs. Mappings allow the template to define a lookup table that associates regions with corresponding AMI IDs, enabling dynamic selection during stack creation across multiple regions. The WGU Cloud Deployment and Operations Study Guide (Section 5.3, CloudFormation Mappings) states, 'The Mappings section enables region-specific configurations, such as mapping AMI IDs to regions (e.g., 'us-east-1': 'ami-123456'), ensuring the correct AMI is used based on the deployment region.' Outputs, WaitCondition, and Resources do not provide this mapping functionality.

Question 9

Question Type: MultipleChoice

(A company is using the CloudWatch agent on its EC2 Linux instances. The company needs to determine how long read requests have waited on the instance's I/O for proper sizing of the instance. Which metric should the company use?)

Options:

A- disk_read_ops
B- disk_io_time
C- diskio_reads
D- diskio_read_bytes

Answer:

B

Explanation:

The `disk_io_time` metric, collected by the CloudWatch agent on EC2 Linux instances, measures the total time (in milliseconds) that read and write requests have waited on I/O, helping to assess disk performance for instance sizing. The WGU Cloud Deployment and Operations Study Guide (Section 4.1, CloudWatch Agent Metrics) states, 'The `disk_io_time` metric, available via the CloudWatch agent, tracks the time spent waiting for I/O operations, providing insight into disk contention for optimizing EC2 instance sizing.' disk_read_ops, diskio_reads, and diskio_read_bytes measure operation counts or bytes, not wait times.

Question 10

Question Type: MultipleChoice

(How are custom metrics grouped in CloudWatch?)

Options:

A- Namespace
B- Service
C- Date

D- Value

Answer:

A

Explanation:

In Amazon CloudWatch, custom metrics are organized and grouped using namespaces. A namespace is a container for CloudWatch metrics that allows you to isolate and categorize metrics from different applications or services. According to the WGU Cloud Deployment and Operations Study Guide (Section 4.1, CloudWatch Metrics), each custom metric must be assigned to a namespace, which acts as a unique identifier to prevent naming collisions and facilitate metric management. Options like Service, Date, and Value are not used for grouping metrics in this context.

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

