

AWS-Certified-Developer-Associate Dumps

Amazon AWS Certified Developer - Associate

<https://www.certleader.com/AWS-Certified-Developer-Associate-dumps.html>



NEW QUESTION 1

You have been doing a lot of testing of your VPC Network by deliberately failing EC2 instances to test whether instances are failing over properly. Your customer who will be paying the AWS bill for all this asks you if he being charged for all these instances. You try to explain to him how the billing works on EC2 instances to the best of your knowledge. What would be an appropriate response to give to the customer in regards to this?

- A. Billing commences when Amazon EC2 AMI instance is completely up and billing ends as soon as the instance starts to shutdown.
- B. Billing commences when Amazon EC2 initiates the boot sequence of an AMI instance and billing ends when the instance shuts down.
- C. Billing only commences only after 1 hour of uptime and billing ends when the instance terminates.
- D. Billing commences when Amazon EC2 initiates the boot sequence of an AMI instance and billing ends as soon as the instance starts to shutdown.

Answer: B

Explanation:

Billing commences when Amazon EC2 initiates the boot sequence of an AMI instance. Billing ends when the instance shuts down, which could occur through a web services command, by running "shutdown -h", or through instance failure.

Reference: <http://aws.amazon.com/ec2/faqs/#Billing>

NEW QUESTION 2

AWS Elastic Load Balancer supports SSL termination.

- A. Tru
- B. For specific availability zones only.
- C. False
- D. Tru
- E. For specific regions only
- F. Tru
- G. For all regions

Answer: D

Explanation:

You can configure your load balancer in ELB (Elastic Load Balancing) to use a SSL certificate in order to improve your system security. The load balancer uses the certificate to terminate and then decrypt requests before sending them to the back-end instances. Elastic Load Balancing uses AWS Identity and Access Management (IAM) to upload your certificate to your load balancer.

Reference: http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/US_SettingUpLoadBalancerHTTPTLS.html

NEW QUESTION 3

Regarding Amazon SNS, when you want to subscribe to a topic and receive notifications to your email, in the Protocol drop-down box, you should select .

- A. Email
- B. Message
- C. SMTP
- D. IMAP

Answer: A

Explanation:

In Amazon SNS, when you want to subscribe to a topic and receive notifications to your email, select Email in the Protocol drop-down box. Enter an email address you can use to receive the notification in the Endpoint field.

Reference: <http://docs.aws.amazon.com/sns/latest/dg/SubscribeTopic.html>

NEW QUESTION 4

is a task coordination and state management service for cloud applications.

- A. Amazon SES
- B. Amazon SWF
- C. Amazon FPS
- D. Amazon SNS

Answer: B

Explanation:

Amazon Simple Workflow (Amazon SWF) is a task coordination and state management service for cloud applications. With Amazon SWF, you can stop writing complex glue-code and state machinery and invest more in the business logic that makes your applications unique.

Reference: <http://aws.amazon.com/swf/>

NEW QUESTION 5

When you create a table with a hash-and-range key, you must define one or more secondary indexes on that table.

- A. False, hash-range key is another name for secondary index
- B. False, it is optional
- C. True
- D. False, when you have Hash-Range key you cannot define Secondary index

Answer: B

Explanation:

When you create a table with a hash-and-range key in DynamoDB, you can also define one or more secondary indexes on that table.
Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/LSI.html>

NEW QUESTION 6

A user is planning to host a scalable dynamic web application on AWS. Which of the services may not be required by the user to achieve automated scalability?

- A. CloudWatch
- B. S3
- C. AutoScaling
- D. AWS EC2 instances

Answer: B

Explanation:

The user can achieve automated scaling by launching different EC2 instances and making them a part of an ELB. Cloudwatch will be used to monitor the resources and based on the scaling need it will trigger policies. AutoScaling is then used to scale up or down the instances.
Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/WhatIsAutoScaling.html>

NEW QUESTION 7

Regarding Amazon SNS, you can send notification messages to mobile devices through any of the following supported push notification services, EXCEPT:

- A. Google Cloud Messaging for Android (GCM)
- B. Apple Push Notification Service (APNS)
- C. Amazon Device Messaging (ADM)
- D. Microsoft Windows Mobile Messaging (MWMM)

Answer: D

Explanation:

In Amazon SNS, you have the ability to send notification messages directly to apps on mobile devices. Notification messages sent to a mobile endpoint can appear in the mobile app as message alerts, badge updates, or even sound alerts. Microsoft Windows Mobile Messaging (MWMM) doesn't exist and is not supported by Amazon SNS.

Reference: <http://docs.aws.amazon.com/sns/latest/dg/SNSMobilePush.html>

NEW QUESTION 8

What kind of service is provided by AWS DynamoDB?

- A. Relational Database
- B. NoSQL Database
- C. Dynamic Database
- D. Document Database

Answer: B

Explanation:

DynamoDB is a fast, fully managed NoSQL database service. Reference: <http://aws.amazon.com/dynamodb/>

NEW QUESTION 9

Doug has created a VPC with CIDR 10.201.0.0/16 in his AWS account. In this VPC he has created a public subnet with CIDR block 10.201.31.0/24. While launching a new EC2 from the console, he is not able to assign the private IP address 10.201.31.6 to this instance. Which is the most likely reason for this issue?

- A. Private IP address 10.201.31.6 is not part of the associated subnet's IP address range.
- B. Private IP address 10.201.31.6 is blocked via ACLs in Amazon infrastructure as a part of platform security.
- C. Private address IP 10.201.31.6 is currently assigned to another interface.
- D. Private IP address 10.201.31.6 is reserved by Amazon for IP networking purpose

Answer: C

Explanation:

In Amazon VPC, you can assign any Private IP address to your instance as long as it is: Part of the associated subnet's IP address range
Not reserved by Amazon for IP networking purposes
Not currently assigned to another interface
Reference: <http://aws.amazon.com/vpc/faqs/>

NEW QUESTION 10

How many types of block devices does Amazon EC2 support?

- A. 5
- B. 1
- C. 2
- D. 4

Answer: C

Explanation:

Amazon EC2 supports 2 types of block devices. Reference:
<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/block-device-mapping-concepts.html>

NEW QUESTION 10

ExamKiller (with AWS account ID H1122223333) has created 50 IAM users for its organization's employees. ExamKiller wants to make the AWS console login URL for all IAM users as: <https://examkiller.signin.aws.amazon.com/console/>. How can this be configured?

- A. Create a bucket with the name ExamKiller and map it with the IAM alias
- B. It is not possible to have capital letters as a part of the alias name
- C. The user needs to use Route 53 to map the ExamKiller domain and IAM URL
- D. For the AWS account, create an alias ExamKiller for the IAM login

Answer: B

Explanation:

If a user wants the URL of the AWS IAM sign-in page to have the company name instead of the AWS account ID, he can create an alias for his AWS account ID. The alias must be unique across all Amazon WebServices products and contain only digits, lowercase letters, and hyphens.

Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/AccountAlias.html>

NEW QUESTION 14

AutoScaling is configured with 3 AZs. Each zone has 5 instances running. If AutoScaling wants to terminate an instance based on the policy action, which instance will it terminate first?

- A. Terminate the first launched instance
- B. Randomly select the instance for termination
- C. Terminate the instance from the AZ which does not have a high AWS load
- D. Terminate the instance from the AZ which has instances running near to the billing hour

Answer: B

Explanation:

Before Auto Scaling selects an instance to terminate, it first identifies the Availability Zone that has more instances than the other Availability Zones used by the group. If all the Availability Zones have the same number of instances, it identifies a random Availability Zone.

Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/us-termination-policy.html>

NEW QUESTION 19

You need to develop and run some new applications on AWS and you know that Elastic Beanstalk and CloudFormation can both help as a deployment mechanism for a broad range of AWS resources. Which of the following statements best describes the differences between Elastic Beanstalk and CloudFormation?

- A. Elastic Beanstalk uses Elastic load balancing and CloudFormation doesn't.
- B. CloudFormation is faster in deploying applications than Elastic Beanstalk.
- C. CloudFormation is much more powerful than Elastic Beanstalk, because you can actually design and script custom resources
- D. Elastic Beanstalk is faster in deploying applications than CloudFormation

Answer: C

Explanation:

These services are designed to complement each other. AWS Elastic Beanstalk provides an environment to easily develop and run applications in the cloud. It is integrated with developer tools and provides a one-stop experience for you to manage the lifecycle of your applications. AWS CloudFormation is a convenient deployment mechanism for a broad range of AWS resources. It supports the infrastructure needs of many different types of applications such as existing enterprise applications, legacy applications, applications built using a variety of AWS resources and container-based solutions (including those built using AWS Elastic Beanstalk).

AWS CloudFormation introduces two new concepts: The template, a JSON-format, text-based file that describes all the AWS resources you need to deploy to run your application and the stack, the set of AWS resources that are created and managed as a single unit when AWS CloudFormation instantiates a template.

Reference: <http://aws.amazon.com/cloudformation/faqs/>

NEW QUESTION 22

Can you SSH to your private machines that reside in a VPC from outside without elastic IP?

- A. Yes, but only if you have direct connect or vpn
- B. Only if you are using a non-US region
- C. Only if you are using a US region
- D. No

Answer: A

Explanation:

The instances that reside in the private subnets of your VPC are not reachable from the Internet, meaning that it is not possible to ssh into them. To interact with them you can use a bastion server, located in a public subnet, that will act as a proxy for them.

You can also connect if you have direct connect or vpn.

Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

NEW QUESTION 24

A user is planning to host a web server as well as an app server on a single EC2 instance which is a part of the public subnet of a VPC. How can the user setup to have two separate public IPs and separate security groups for both the application as well as the web server?

- A. Launch a VPC instance with two network interface
- B. Assign a separate security group to each and AWS will assign a separate public IP to them.
- C. Launch VPC with two separate subnets and make the instance a part of both the subnets.
- D. Launch a VPC instance with two network interface

- E. Assign a separate security group and elastic IP to them.
- F. Launch a VPC with ELB such that it redirects requests to separate VPC instances of the public subne

Answer: C

Explanation:

If you need to host multiple websites(with different IPs) on a single EC2 instance, the following is the suggested method from AWS.

Launch a VPC instance with two network interfaces

Assign elastic IPs from VPC EIP pool to those interfaces (Because, when the user has attached more than one network interface with an instance, AWS cannot assign public IPs to them.)

Assign separate Security Groups if separate Security Groups are needed

This scenario also helps for operating network appliances, such as firewalls or load balancers that have multiple private IP addresses for each network interface.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/MultipleIP.html>

NEW QUESTION 28

An online gaming site asked you if you can deploy a database that is a fast, highly scalable NoSQL database service in AWS for a new site that he wants to build. Which database should you recommend?

- A. Amazon Redshift
- B. Amazon SimpleDB
- C. Amazon DynamoDB
- D. Amazon RDS

Answer: C

Explanation:

Amazon DynamoDB is ideal for database applications that require very low latency and predictable performance at any scale but don't need complex querying capabilities like joins or transactions. Amazon DynamoDB is a fully-managed NoSQL database service that offers high performance, predictable throughput and low cost. It is easy to set up, operate, and scale.

With Amazon DynamoDB, you can start small, specify the throughput and storage you need, and easily scale your capacity requirements on the fly. Amazon DynamoDB automatically partitions data over a number of servers to meet your request capacity. In addition, DynamoDB automatically replicates your data synchronously across multiple Availability Zones within an AWS Region to ensure high-availability and data durability.

Reference: https://aws.amazon.com/running_databases/#dynamodb_anchor

NEW QUESTION 32

How long are the messages kept on an SQS queue by default?

- A. If a message is not read, it is never deleted
- B. 2 weeks
- C. 1 day
- D. 4 days

Answer: D

Explanation:

The SQS message retention period is configurable and can be set anywhere from 1 minute to 2 weeks. The default is 4 days and once the message retention limit is reached your messages will be automatically deleted. The option for longer message retention provides greater flexibility to allow for longer intervals between message production and consumption.

Reference: <https://aws.amazon.com/sqs/faqs/>

NEW QUESTION 33

A user has setup an application on EC2 which uses the IAM user access key and secret access key to make secure calls to S3. The user wants to temporarily stop the access to S3 for that IAM user. What should the root owner do?

- A. Delete the IAM user
- B. Change the access key and secret access key for the users
- C. Disable the access keys for the IAM user
- D. Stop the instance

Answer: C

Explanation:

If the user wants to temporarily stop the access to S3 the best solution is to disable the keys. Deleting the user will result in a loss of all the credentials and the app will not be useful in the future. If the user stops the instance IAM users can still access S3. The change of the key does not help either as they are still active. The best possible solution is to disable the keys.

Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/ManagingCredentials.html>

NEW QUESTION 37

A user wants to configure AutoScaling which scales up when the CPU utilization is above 70% and scales down when the CPU utilization is below 30%. How can the user configure AutoScaling for the above mentioned condition?

- A. Use AutoScaling with a schedule
- B. Configure ELB to notify AutoScaling on load increase or decrease
- C. Use dynamic AutoScaling with a policy
- D. Use AutoScaling by manually modifying the desired capacity during a condition

Answer: C

Explanation:

The user can configure the AutoScaling group to automatically scale up and then scale down based on the specified conditions. To configure this, the user must setup policies which will get triggered by the CloudWatch alarms.

Reference:

<http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-scale-based-on-demand.html>

NEW QUESTION 38

A user has created an application which sends data to a log file. The server hosting the log files can be unavailable due to any reason. The user wants to make it so that whenever the log server is up it should be receiving the messages. Which of the below mentioned AWS services helps achieve this functionality?

- A. AWS Simple Workflow
- B. AWS Simple Task Service
- C. AWS Simple Notification Service
- D. AWS Simple Queue Service

Answer: D

Explanation:

Amazon Simple Queue Service (SQS) is a fast, reliable, scalable, and fully managed message queuing service. SQS provides a simple and cost-effective way to decouple the components of an application. The user can use SQS to transmit any volume of data without losing messages or requiring other services to always be available. Using SQS, the application has to just send the data to SQS and SQS transmits it to the log file whenever it is available.

Reference: <http://aws.amazon.com/sqs/>

NEW QUESTION 41

A user is creating an ELB with VPC. Which of the following options is available as a part of the "Add EC2 instances" page?

- A. Select Subnet
- B. Select IAM
- C. Select ENI
- D. Select VPC

Answer: A

Explanation:

When a user is launching an ELB with VPC, he/she has to select the options, such as subnet and security group before selecting the instances part of that subnet.

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-getting-started.html>

NEW QUESTION 45

What is the maximum time messages can be stored in SQS?

- A. 14 days
- B. one month
- C. 4 days
- D. 7 days

Answer: A

Explanation:

A message can be stored in the Simple Queue Service (SQS) from 1 minute up to a maximum of 14 days. Reference:

http://aws.amazon.com/sqs/faqs/#How_long_can_I_keep_my_messages_in_Amazon_SQS_queues

NEW QUESTION 49

A user is launching an AWS RDS instance with MySQL. The user wants to enable the Multi AZ feature. Which of the below mentioned parameters will not be allowed to configure by RDS?

- A. Availability Zone
- B. Region
- C. DB subnet group
- D. Database port

Answer: A

Explanation:

If the user is launching RDS with Multi AZ the user cannot provision the Availability Zone. RDS is launched automatically instead

Reference: <https://console.aws.amazon.com/rds/>

NEW QUESTION 54

A user has launched an EBS backed Linux instance. How can a user detach the root device and attach it to another instance as a secondary volume?

- A. Unmount the root volume first and then detach it
- B. It is not possible to mount the root volume to some other instance
- C. Stop the first instance and then attach instance's root volume as a new volume to the other instance
- D. It is not possible to mount the root device as a secondary volume on the other instance

Answer: C

Explanation:

If an Amazon EBS volume is the root device of an instance, it cannot be detached unless the instance is in the stopped state.
Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-detaching-volume.html>

NEW QUESTION 55

When using Amazon SQS how much data can you store in a message?

- A. 8 KB
- B. 2 KB
- C. 16 KB
- D. 4 KB

Answer: A

Explanation:

With Amazon SQS version 2008-01-01, the maximum message size for both SOAP and Query requests is 8KB.
If you need to send messages to the queue that are larger than 8 KB, AWS recommends that you split the information into separate messages. Alternatively, you could use Amazon S3 or Amazon SimpleDB to hold the information and include the pointer to that information in the Amazon SQS message.
If you send a message that is larger than 8KB to the queue, you will receive a MessageTooLong error with HTTP code 400.
Reference: <https://aws.amazon.com/items/1343?externalID=1343>

NEW QUESTION 58

In regard to DynamoDB, what is the Global secondary index?

- A. An index with a hash and range key that can be different from those on the table.
- B. An index that has the same range key as the table, but a different hash key
- C. An index that has the same hash key and range key as the table
- D. An index that has the same hash key as the table, but a different range key

Answer: A

Explanation:

Global secondary index - an index with a hash and range key that can be different from those on the table.
Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DataModel.html>

NEW QUESTION 62

You have been given a scope to deploy some AWS infrastructure for a large organization. The requirements are that you will have a lot of EC2 instances but may need to add more when the average utilization of your Amazon EC2 fleet is high and conversely remove them when CPU utilization is low. Which AWS services would be best to use to accomplish this?

- A. Amazon CloudFront, Amazon CloudWatch and Elastic Load Balancing.
- B. Auto Scaling, Amazon CloudWatch and AWS CloudTrail.
- C. Auto Scaling, Amazon CloudWatch and Elastic Load Balancing.
- D. Auto Scaling, Amazon CloudWatch and AWS Elastic Beanstalk

Answer: C

Explanation:

Auto Scaling enables you to follow the demand curve for your applications closely, reducing the need to manually provision Amazon EC2 capacity in advance. For example, you can set a condition to add new Amazon EC2 instances in increments to the Auto Scaling group when the average utilization of your Amazon EC2 fleet is high; and similarly, you can set a condition to remove instances in the same increments when CPU utilization is low. If you have predictable load changes, you can set a schedule through Auto Scaling to plan your scaling activities. You can use Amazon CloudWatch to send alarms to trigger scaling activities and Elastic Load Balancing to help distribute traffic to your instances within Auto Scaling groups. Auto Scaling enables you to run your Amazon EC2 fleet at optimal utilization. Reference: <http://aws.amazon.com/autoscaling/>

NEW QUESTION 63

You are building an online store on AWS that uses SQS to process your customer orders. Your backend system needs those messages in the same sequence the customer orders have been put in. How can you achieve that?

- A. You can do this with SQS but you also need to use SWF
- B. Messages will arrive in the same order by default
- C. You can use sequencing information on each message
- D. It is not possible to do this with SQS

Answer: C

Explanation:

Amazon SQS is engineered to always be available and deliver messages. One of the resulting tradeoffs is that SQS does not guarantee first in, first out delivery of messages. For many distributed applications, each message can stand on its own, and as long as all messages are delivered, the order is not important. If your system requires that order be preserved, you can place sequencing information in each message, so that you can reorder the messages when the queue returns them. Reference: <http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/welcome.html>

NEW QUESTION 66

A user has launched an EC2 instance and installed a website with the Apache webserver. The webserver is running but the user is not able to access the website from the internet. What can be the possible reason for this failure?

- A. The security group of the instance is not configured properly.
- B. The instance is not configured with the proper key-pairs.
- C. The Apache website cannot be accessed from the internet.
- D. Instance is not configured with an elastic I

Answer: A

Explanation:

In Amazon Web Services, when a user has configured an instance with Apache, the user needs to ensure that the ports in the security group are opened as configured in Apache config. E.g. If Apache is running on port 80, the user should open port 80 in the security group.
Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html>

NEW QUESTION 71

When you use the AWS Elastic Beanstalk console to deploy a new application you'll need to upload a source bundle and it should .

- A. Consist of a single .zip file
- B. Consist of a single .war file
- C. Consist of a single .zip file or .war file
- D. Consist of a folder with all files

Answer: C

Explanation:

When you use the AWS Elastic Beanstalk console to deploy a new application or an application version, you'll need to upload a source bundle. Your source bundle must meet the following requirements: Consist of a single .zip file or .war file
Not exceed 512 MB
Not include a parent folder or top-level directory (subdirectories are fine) Reference:
<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.deployment.source.html>

NEW QUESTION 73

How does Amazon SQS allow multiple readers to access the same message queue without losing messages or processing them many times?

- A. By identifying a user by his unique id
- B. By using unique cryptography
- C. Amazon SQS queue has a configurable visibility timeout.
- D. Multiple readers can't access the same message queue

Answer: C

Explanation:

Every Amazon SQS queue has a configurable visibility timeout. For the designated amount of time after a message is read from a queue, it will not be visible to any other reader. As long as the amount of time that it takes to process the message is less than the visibility timeout, every message will be processed and deleted. In the event that the component processing the message fails or becomes unavailable, the message will again become visible to any component reading the queue once the visibility timeout ends. This allows you to have many components all reading messages from the same queue, with each working to process different messages.
Reference: <https://aws.amazon.com/sqs/faqs/>

NEW QUESTION 77

Which of the following device names is recommended for an EBS volume that can be attached to an Amazon EC2 Instance running Windows?

- A. xvd[a-e]
- B. /mnt/sd[b-e]
- C. xvd[f-p]
- D. /dev/sda1

Answer: C

Explanation:

The xvd[f-p] is the recommended device name for EBS volumes that can be attached to the Amazon EC2 Instances running on Windows.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/device_naming.html

NEW QUESTION 81

A user is trying to create a list of IAM users with the AWS console. When the IAM users are created which of the below mentioned credentials will be enabled by default for the user?

- A. IAM access key and secret access key
- B. IAM X.509 certificates
- C. Nothin
- D. Everything is disabled by default
- E. IAM passwords

Answer: C

Explanation:

Newly created IAM users have no password and no access key (access key ID and secret access key). If the user needs to administer your AWS resources using the AWS Management Console, you can create a password for the user. If the user needs to interact with AWS programmatically (using the command line interface (CLI), the AWS SDK, or service-specific APIs), you can create an access key for that user. The credentials you create for users are what they use to uniquely identify themselves to AWS.

Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_WorkingWithGroupsAndUsers.html

NEW QUESTION 86

A user has created photo editing software and hosted it on EC2. The software accepts requests from the user about the photo format and resolution and sends a message to S3 to enhance the picture accordingly. Which of the below mentioned AWS services will help make a scalable software with the AWS infrastructure in this scenario?

- A. AWS Elastic Transcoder
- B. AWS Simple Notification Service
- C. AWS Simple Queue Service
- D. AWS Glacier

Answer: C

Explanation:

Amazon Simple Queue Service (SQS) is a fast, reliable, scalable, and fully managed message queuing service. SQS provides a simple and cost-effective way to decouple the components of an application. The user can configure SQS, which will decouple the call between the EC2 application and S3. Thus, the application does not keep waiting for S3 to provide the data.

Reference: <http://aws.amazon.com/sqs/faqs/>

NEW QUESTION 90

A user has created a blank EBS volume in the US-East-1 region. The user is unable to attach the volume to a running instance in the same region. What could be the possible reason for this?

- A. The instance must be in a running state
- B. It is required to stop the instance to attach volume
- C. The AZ for the instance and volume are different
- D. The instance is from an instance store backed AMI
- E. The instance has enabled the volume attach protection

Answer: B

Explanation:

An EBS volume provides persistent data storage. The user can attach a volume to any instance provided they are both in the same AZ. Even if they are in the same region but in a different AZ, it will not be able to attach the volume to that instance.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html>

NEW QUESTION 92

A root account owner is trying to setup an additional level of security for all his IAM users. Which of the below mentioned options is a recommended solution for the account owner?

- A. Enable access key and secret access key for all the IAM users
- B. Enable MFA for all IAM users
- C. Enable the password for all the IAM users
- D. Enable MFA for the root account

Answer: B

Explanation:

Multi-Factor Authentication adds an extra level of security for all the users. The user can enable MFA for all IAM users which ensures that each user has to provide an extra six digit code for authentication. Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_ManagingMFA.html

NEW QUESTION 97

A user is creating multiple IAM users. What advice should be given to him to enhance the security?

- A. Grant least privileges to the individual user
- B. Grant all higher privileges to the group
- C. Grant less privileges for user, but higher privileges for the group
- D. Grant more privileges to the user, but least privileges to the group

Answer: A

Explanation:

It is a recommended rule that the root user should grant the least privileges to the IAM user or the group. The higher the privileges, the more problems it can create.

Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/IAMBestPractices.html>

NEW QUESTION 98

A user has set an IAM policy where it allows all requests if a request from IP 10.10.10.1/32. Another policy allows all the requests between 5 PM to 7 PM. What will happen when a user is requesting access from IP 10.10.10.1/32 at 6 PM?

- A. IAM will throw an error for policy conflict
- B. It is not possible to set a policy based on the time or IP
- C. It will deny access
- D. It will allow access

Answer: D

Explanation:

With regard to IAM, when a request is made, the AWS service decides whether a given request should be allowed or denied. The evaluation logic follows these rules:

By default, all requests are denied. (In general, requests made using the account credentials for resources in the account are always allowed.)

An explicit allow policy overrides this default. An explicit deny policy overrides any allows. Reference:

http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_EvaluationLogic.html

NEW QUESTION 101

Is it possible to create an S3 bucket accessible only by a certain IAM user, using policies in a CloudFormation template?

- A. No, you can only create the S3 bucket but not the IAM user.
- B. S3 is not supported by CloudFormation.
- C. Yes, all these resources can be created using a CloudFormation template
- D. No, in the same template you can only create the S3 bucket and the relative policy

Answer: C

Explanation:

With AWS Identity and Access Management (IAM), you can create IAM users to control who has access to which resources in your AWS account. You can use IAM with AWS CloudFormation to control what AWS CloudFormation actions users can perform, such as view stack templates, create stacks, or delete stacks. In addition to AWS CloudFormation actions, you can manage what AWS services and resources are available to each user.

NEW QUESTION 105

A user has launched an RDS instance. The user has created 3 databases on the same server. What can the maximum size be for each database?

- A. The size of each DB cannot be more than 3 TB
- B. It is not possible to have more than one DB on a single instance
- C. The total instance storage size cannot be more than 3 TB
- D. The size of each DB cannot be more than 1 TB

Answer: C

Explanation:

The AWS RDS DB instance is an isolated DB environment provided by AWS in which the user can create more than 1 database. The maximum size of the instance should be between 5 GB and 3 TB. The size of each DB can be anything in this range.

Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html>

NEW QUESTION 107

A user has created an RDS instance with MySQL. The user is using the HeidiSQL client to connect with the RDS DB. The client is unable to connect to DB from his home machine. What is a possible reason for the failure?

- A. The user has to open port 80 in the RDS security group to connect with RDS DNS
- B. The security group is not configured to allow a request from the user's IP on port 3306
- C. You can never connect to RDS from your desktop
- D. The user has to open port 22 in the RDS security group to connect with RDS DNS

Answer: B

Explanation:

If the user needs to connect to RDS then he has to open port 3306 in the RDS security group for his IP address.

Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html>

NEW QUESTION 112

A user is creating a new EBS volume from an existing snapshot. The snapshot size shows 10 GB. Can the user create a volume of 30 GB from that snapshot?

- A. Provided the original volume has set the change size attribute to true
- B. Yes
- C. Provided the snapshot has the modify size attribute set as true
- D. No

Answer: B

Explanation:

A user can always create a new EBS volume of a higher size than the original snapshot size. The user cannot create a volume of a lower size. When the new volume is created the size in the instance will be shown as the original size. The user needs to change the size of the device with `resize2fs` or other OS specific commands.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-expand-volume.html>

NEW QUESTION 115

A user has launched a MySQL RDS. The user wants to plan for the DR and automate the snapshot. Which of the below mentioned functionality offers this option with RDS?

- A. Copy snapshot
- B. Automated synchronization
- C. Snapshot
- D. Automated backup

Answer: D

Explanation:

Amazon RDS provides two different methods for backing up and restoring the Amazon DB instances: automated backups and DB snapshots. Automated backups automatically back up the DB instance during a specific, user-definable backup window, and keep the backups for a limited, user-specified period of time. Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.BackingUpAndRestoringAmazonRDSInstances.html>

NEW QUESTION 120

When AutoScaling is launching a new instance based on condition, which of the below mentioned policies will it follow?

- A. Based on the criteria defined with cross zone Load balancing
- B. Launch an instance which has the highest load distribution
- C. Launch an instance in the AZ with the fewest instances
- D. Launch an instance in the AZ which has the highest instances

Answer: C

Explanation:

AutoScaling attempts to distribute instances evenly between the Availability Zones that are enabled for the user's AutoScaling group. Auto Scaling does this by attempting to launch new instances in the Availability Zone with the fewest instances.

Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html

NEW QUESTION 124

Which of the below mentioned options can be a good use case for storing content in AWS RRS?

- A. Storing mission critical data Files
- B. Storing infrequently used log files
- C. Storing a video file which is not reproducible
- D. Storing image thumbnails

Answer: D

Explanation:

AWS RRS provides the same functionality as AWS S3, but at a cheaper rate. It is ideally suited for non-mission, critical applications, such as files which can be reproduced.

Reference: <http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingRRS.html>

NEW QUESTION 129

A user is using an EBS backed instance. Which of the below mentioned statements is true?

- A. The user will be charged for volume and instance only when the instance is running
- B. The user will be charged for the volume even if the instance is stopped
- C. The user will be charged only for the instance running cost
- D. The user will not be charged for the volume if the instance is stopped

Answer: B

Explanation:

If a user has launched an EBS backed instance, the user will be charged for the EBS volume even though the instance is in a stopped state. The instance will be charged for the EC2 hourly cost only when it is running.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-detaching-volume.html>

NEW QUESTION 134

A user is trying to create a policy for an IAM user from the AWS console. Which of the below mentioned options is not available to the user while configuring policy?

- A. Use policy generator to create policy
- B. Use custom policy to create policy
- C. Use policy simulator to create policy
- D. Assign No permission

Answer: C

Explanation:

When a user is trying to create a policy from the AWS console, it will have options such as create policy from templates or use a policy generator. The user can also define a custom policy or chose the option to have no permission. The policy simulator is not available in the console.

Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/IAMBestPractices.html>

NEW QUESTION 135

A user is planning to use EBS for his DB requirement. The user already has an EC2 instance running in the VPC private subnet. How can the user attach the EBS volume to a running instance?

- A. The user must create EBS within the same VPC and then attach it to a running instance.
- B. The user can create EBS in the same zone as the subnet of instance and attach that EBS to instance.
- C. It is not possible to attach an EBS to an instance running in VPC until the instance is stopped.
- D. The user can specify the same subnet while creating EBS and then attach it to a running instance

Answer: B

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. The VPC is always specific to a region. The user can create a VPC which can span multiple Availability Zones by adding one or more subnets in each Availability Zone. The instance launched will always be in the same availability zone of the respective subnet. When creating an EBS the user cannot specify the subnet or VPC. However, the user must create the EBS in the same zone as the instance so that it can attach the EBS volume to the running instance.

Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.htm#VPCSubnet

NEW QUESTION 139

A user is trying to find the state of an S3 bucket with respect to versioning. Which of the below mentioned states AWS will not return when queried?

- A. versioning-enabled
- B. versioning-suspended
- C. unversioned
- D. versioned

Answer: D

Explanation:

S3 buckets can be in one of the three states: unversioned (the default), versioning-enabled or versioning-suspended. The bucket owner can configure the versioning state of a bucket. The versioning state applies to all (never some) of the objects in that bucket. The first time owner enables a bucket for versioning, objects in it are thereafter always versioned and given a unique version ID.

Reference: <http://docs.aws.amazon.com/AmazonS3/latest/dev/Versioning.html>

NEW QUESTION 141

How do you configure SQS to support longer message retention?

- A. Set the IVmessageRetentionPeriod attribute using the SetQueueAttributes method
- B. Using a Lambda function
- C. You can't
- D. It is set to 14 days and cannot be changed
- E. You need to request it from AWS

Answer: A

Explanation:

To configure the message retention period, set the IVmessageRetentionPeriod attribute using the SetQueueAttributes method. This attribute is used to specify the number of seconds a message will be retained by SQS. Currently the default value for the message retention period is 4 days. Using the IVmessageRetentionPeriod attribute, the message retention period can be set anywhere from 60 seconds (1 minute), up to 1209600 seconds (14 days).

Reference: <https://aws.amazon.com/sqs/faqs/>

NEW QUESTION 144

A user has enabled serverside encryption with S3. The user downloads the encrypted object from S3. How can the user decrypt it?

- A. S3 does not support server side encryption
- B. S3 provides a server side key to decrypt the object
- C. The user needs to decrypt the object using their own private key
- D. S3 manages encryption and decryption automatically

Answer: D

Explanation:

If the user is using the server-side encryption feature, Amazon S3 encrypts the object data before saving it on disks in its data centres and decrypts it when the user downloads the objects. Thus, the user is free from the tasks of managing encryption, encryption keys, and related tools.

Reference: <http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingEncryption.html>

NEW QUESTION 146

A user is trying to configure access with S3. Which of the following options is not possible to provide access to the S3 bucket / object?

- A. Define the policy for the IAM user
- B. Define the ACL for the object
- C. Define the policy for the object
- D. Define the policy for the bucket

Answer: C

Explanation:

Amazon S3 offers access policy options broadly categorized as resource-based policies and user policies.

Access policies, such as ACL and resource policy can be attached to the bucket. With the object the user can only have ACL and not an object policy. The user can also attach access policies to the IAM users in the account. These are called user policies.

Reference: <http://docs.aws.amazon.com/AmazonS3/latest/dev/s3-access-control.html>

NEW QUESTION 149

ExamKiller (with AWS account ID 111122223333) has created 50 IAM users for its orgAMzation's employees. What will be the AWS console URL for these associates?

- A. [https:// 111122223333.signin.aws.amazon.com/console/](https://111122223333.signin.aws.amazon.com/console/)
- B. [https:// signin.aws.amazon.com/console/](https://signin.aws.amazon.com/console/)
- C. <https://signin.aws.amazon.com/111122223333/console/>
- D. <https://signin.aws.amazon.com/console/111122223333/>

Answer: A

Explanation:

When an orgAMzation is using AWS IAM for creating various users and manage their access rights, the IAM user cannot use the login URL <http://aws.amazon.com/console> to access AWS management console. The console login URL for the IAM user will have AWS account ID of that orgAMzation to identify the IAM user belongs to particular account. The AWS console login URL for the IAM user will be [https:// <AWS_Account_ID>.signin.aws.amazon.com/console/](https://<AWS_Account_ID>.signin.aws.amazon.com/console/). In this case it will be [https:// 111122223333.signin.aws.amazon.com/console/](https://111122223333.signin.aws.amazon.com/console/)
Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/AccountAlias.html>

NEW QUESTION 151

True or False: AWS CloudFormation allows you to create Microsoft Windows stacks.

- A. False, AWS CloudFormation does not support Microsoft Windows.
- B. False, Amazon doesn't support Microsoft Windows.
- C. False, you cannot create Windows stacks.
- D. True

Answer: D

Explanation:

AWS CloudFormation allows you to create Microsoft Windows stacks based on Amazon EC2 Windows Amazon Machine Images (AMIs) and provides you with the ability to install software, to use remote desktop to access your stack, and to update and configure your stack.
Reference: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-windows-stacks.html>

NEW QUESTION 155

Which of the following solutions is not supported by DynamoDB:

- A. Hash secondary index
- B. Local secondary index
- C. Hash Primary Key
- D. Global secondary index

Answer: A

Explanation:

In DynamoDB, a secondary index is a data structure that contains a subset of attributes from a table, along with an alternate key to support Query operations. DynamoDB supports the following two types of secondary indexes:
Local secondary index is an index that has the same hash key as the table, but a different range key. A local secondary index is "local" in the sense that every partition of a local secondary index is scoped to a table partition that has the same hash key.
Global secondary index is an index with a hash and range key that can be different from those on the table. A global secondary index is considered "global" because queries on the index can span all of the data in a table, across all partitions.
Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DataModel.html>

NEW QUESTION 158

An ELB is diverting traffic across 5 instances. One of the instances was unhealthy only for 20 minutes. What will happen after 20 minutes when the instance becomes healthy?

- A. ELB will never divert traffic back to the same instance
- B. ELB will not automatically send traffic to the same instanc
- C. However, the user can configure to start sending traffic to the same instance
- D. ELB starts sending traffic to the instance once it is healthy
- E. ELB terminates the instance once it is unhealth
- F. Thus, the instance cannot be healthy after 10 minutes

Answer: C

Explanation:

AWS Elastic Load Balancing continuously checks the health of an instance. If one of the instances is unhealthy it stops sending traffic to it and automatically reroutes the traffic to the remaining running EC2 instances. If the failed EC2 instance is restored, Elastic Load Balancing will again start sending traffic to that instance.
Reference: <http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/SvcIntro.htm>

NEW QUESTION 160

Which OS does the current version of AWS Elastic Beanstalk use?

- A. Amazon Linux AMI, Windows Server 2003 R2 AMI or the Windows Server 2008 R2 AMI
- B. Amazon Linux AMI only
- C. Amazon Linux AMI or the Windows Server 2008 R2 AMI
- D. Windows Sewer 2008 R2 AMI only

Answer: C

Explanation:

The current version of AWS Elastic Beanstalk uses the Amazon Linux AMI or the Windows Server 2008 R2 AMI.

Reference: <https://aws.amazon.com/elasticbeanstalk/faqs/>

NEW QUESTION 161

A user is creating an EBS volume. He asks for your advice. Which advice mentioned below should you not give to the user for creating an EBS volume?

- A. Take the snapshot of the volume when the instance is stopped
- B. Stripe multiple volumes attached to the same instance
- C. Create an AMI from the attached volume
- D. Attach multiple volumes to the same instance

Answer: C

Explanation:

When a user creates an EBS volume, the user can attach it to a running instance. The user can attach multiple volumes to the same instance and stripe them together to increase the I/O. The user can take a snapshot from the existing volume but cannot create an AMI from the volume. However, the user can create an AMI from a snapshot.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumes.html>

NEW QUESTION 163

Your supervisor has asked you to build a simple file synchronization service for your department. He doesn't want to spend too much money and he wants to be notified of any changes to files by email. What do you think would be the best Amazon service to use for the email solution?

- A. Amazon CloudSearch
- B. Amazon Elastic Transcoder
- C. Amazon SES
- D. Amazon AppStream

Answer: C

Explanation:

File change notifications can be sent via email to users following the resource with Amazon Simple Email Service (Amazon SES), an easy-to-use, cost-effective email solution.

Reference: http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_filesync_08.pdf

NEW QUESTION 165

An organization has 10 departments. The organization wants to track the AWS usage of each department. Which of the below mentioned options meets the requirement?

- A. Setup IAM groups for each department and track their usage
- B. Create separate accounts for each department, but use consolidated billing for payment and tracking
- C. Create separate accounts for each department and track them separately
- D. Setup IAM users for each department and track their usage

Answer: B

Explanation:

The cost of an IAM user or groups can never be tracked separately for the purpose of billing. The best solution in this case is to create a separate account for each department and use consolidated billing. Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/IAM_introduction.html

NEW QUESTION 169

What is the maximum size for messages stored in SQS?

- A. 256KB
- B. 128KB
- C. 1024KB
- D. 64KB

Answer: A

Explanation:

By default, SQS queues allow you to send the largest supported payload size, currently 256KB. You can choose to specify a limit on how many bytes can be sent per payload, using the `MaximumMessageSize` attribute of the `SetQueueAttributes` method.

Reference: <http://aws.amazon.com/sqs/faqs/>

NEW QUESTION 172

A user has configured an automated backup between 5 AM — 5:30 AM for the MySQL RDS DB. Will the performance of RDS get frozen momentarily during a backup?

- A. No
- B. Yes, only if the instance size is smaller than large size
- C. Yes, provided it is a single zone implementation
- D. Yes, always

Answer: C

Explanation:

Amazon RDS provides two different methods for backing up and restoring the Amazon DB instances. A brief I/O freeze, typically lasting a few seconds, occurs during both automated backups and DB snapshot operations on Single-AZ DB instances.

Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.BackingUpAndRestoringAmazonRDSInstances.html>

NEW QUESTION 177

To scale up the AWS resources using manual AutoScaling, which of the below mentioned parameters should the user change?

- A. Maximum capacity
- B. Desired capacity
- C. Preferred capacity
- D. Current capacity

Answer: B

Explanation:

The Manual Scaling as part of Auto Scaling allows the user to change the capacity of Auto Scaling group. The user can add / remove EC2 instances on the fly. To execute manual scaling, the user should modify the desired capacity. AutoScaling will adjust instances as per the requirements. If the user is trying to CLI, he can use command `as-set-desired-capacity <Auto Scaling Group Name> --desired-capacity <New Capacity>`

Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-manual-scaling.html>

NEW QUESTION 181

Does DynamoDB support in-place atomic updates?

- A. It is not defined
- B. Yes
- C. It does support in-place non-atomic updates
- D. No

Answer: B

Explanation:

DynamoDB supports in-place atomic updates. Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/WorkingWithItems.html#WorkingWithItems.AtomicCounters>

NEW QUESTION 183

A user is having access to objects of an S3 bucket which is not owned by him. If he is trying to set the objects of that bucket public, which of the below mentioned options may be a right fit for this action?

- A. Make the bucket public with full access
- B. Define the policy for the bucket
- C. Provide ACL on the object
- D. Create an IAM user with permission

Answer: C

Explanation:

An S3 object ACL is the only way to manage access to objects which are not owned by the bucket owner. An AWS account that owns the bucket can grant another AWS account permission to upload objects. The bucket owner does not own these objects. The AWS account that created the object must grant permissions using object ACLs.

Reference: <http://docs.aws.amazon.com/AmazonS3/latest/dev/access-policy-alternatives-guidelines.html>

NEW QUESTION 185

A user has created an EBS volume with 1000 IOPS. What is the average IOPS that the user will get for most of the year as per EC2 SLA if the instance is attached to the EBS optimized instance?

- A. 900
- B. 990
- C. 950
- D. 1000

Answer: A

Explanation:

As per AWS SLA if the instance is attached to an EBS-Optimized instance, then the Provisioned IOPS volumes are designed to deliver within 10% of the provisioned IOPS performance 99.9% of the time in a given year. Thus, if the user has created a volume of 1000 IOPS, the user will get a minimum 900 IOPS 99.9% time of the year.

Reference: <http://aws.amazon.com/ec2/faqs/>

NEW QUESTION 188

Which of the following is an example of a good DynamoDB hash key schema for provisioned throughput efficiency?

- A. User ID, where the application has many different users.
- B. Status Code where most status codes are the same
- C. Device ID, where one is by far more popular than all the others.
- D. Game Type, where there are three possible game types

Answer: A

NEW QUESTION 190

Which of the following are correct statements with policy evaluation logic in AWS Identity and Access Management? Choose 2 answers

- A. By default, all requests are denied
- B. An explicit allow overrides an explicit deny
- C. An explicit allow overrides default deny.
- D. An explicit deny does not override an explicit allow
- E. By default, all request are allowed

Answer: AC

NEW QUESTION 193

Company C has recently launched an online commerce site for bicycles on AWS. They have a "Product" DynamoDB table that stores details for each bicycle, such as, manufacturer, color, price, quantity and size to display in the online store. Due to customer demand, they want to include an image for each bicycle along with the existing details.

Which approach below provides the least impact to provisioned throughput on the "Product" table?

- A. Serialize the image and store it in multiple DynamoDB tables
- B. Create an "Images" DynamoDB table to store the Image with a foreign key constraint to the "Product" table
- C. Add an image data type to the "Product" table to store the images in binary format
- D. Store the images in Amazon S3 and add an S3 URL pointer to the "Product" table item for each image

Answer: D

NEW QUESTION 198

Company D is running their corporate website on Amazon S3 accessed from <http://www.companyd.com>. Their marketing team has published new web fonts to a separate S3 bucket accessed by the S3 endpoint <https://s3-us-west1.amazonaws.com/cdfonts>. While testing the new web fonts, Company D recognized the web fonts are being blocked by the browser. What should Company D do to prevent the web fonts from being blocked by the browser?

- A. Enable versioning on the cdfonts bucket for each web font
- B. Create a policy on the cdfonts bucket to enable access to everyone
- C. Add the Content-NI D5 header to the request for webfonts in the cdfonts bucket from the website
- D. Configure the cdfonts bucket to allow cross-origin requests by creating a CORS configuration

Answer: D

NEW QUESTION 201

When uploading an object, what request header can be explicitly specified in a request to Amazon S3 to encrypt object data when saved on the server side?

- A. x-amz-storage-class
- B. Content-MD5
- C. x-amz-security-token
- D. x-amz-server-side-encryption

Answer: D

NEW QUESTION 202

In AWS, which security aspects are the customer's responsibility? Choose 4 answers

- A. Life-cycle management of IAM credentials
- B. Decommissioning storage devices
- C. Security Group and ACL (Access Control List) settings
- D. Encryption of EBS (Elastic Block Storage) volumes
- E. Controlling physical access to compute resources
- F. Patch management on the EC2 instance's operating system

Answer: ABCF

NEW QUESTION 206

What item operation allows the retrieval of multiple items from a DynamoDB table in a single API call?

- A. GetItem
- B. BatchGetItem
- C. GetMultipleItems
- D. GetItemRange

Answer: B

NEW QUESTION 211

An application stores payroll information nightly in DynamoDB for a large number of employees across hundreds of offices. Item attributes consist of individual name, office identifier, and cumulative daily hours. Managers run reports for ranges of names working in their office. One query is, "Return all items in this office for names starting with A through E". Which table configuration will result in the lowest impact on provisioned throughput for this query?

- A. Configure the table to have a hash index on the name attribute, and a range index on the office identifier
- B. Configure the table to have a range index on the name attribute, and a hash index on the office identifier
- C. Configure a hash index on the name attribute and no range index

D. Configure a hash index on the office Identifier attribute and no range index

Answer: B

NEW QUESTION 212

When a Simple Queue Service message triggers a task that takes 5 minutes to complete, which process below will result in successful processing of the message and remove it from the queue while minimizing the chances of duplicate processing?

- A. Retrieve the message with an increased visibility timeout, process the message, delete the message from the queue
- B. Retrieve the message with an increased visibility timeout, delete the message from the queue, process the message
- C. Retrieve the message with increased DelaySeconds, process the message, delete the message from the queue
- D. Retrieve the message with increased DelaySeconds, delete the message from the queue, process the message

Answer: A

NEW QUESTION 215

Which of the following services are key/value stores? Choose 3 answers

- A. Amazon ElastiCache
- B. Simple Notification Service
- C. DynamoDB
- D. Simple Workflow Service
- E. Simple Storage Service

Answer: ABC

NEW QUESTION 219

Company A has an S3 bucket containing premier content that they intend to make available to only paid subscribers of their website. The S3 bucket currently has default permissions of all objects being private to prevent inadvertent exposure of the premier content to non-paying website visitors. How can Company A provide only paid subscribers the ability to download a premier content file in the S3 bucket?

- A. Apply a bucket policy that grants anonymous users to download the content from the S3 bucket
- B. Generate a pre-signed object URL for the premier content file when a paid subscriber requests a download
- C. Add a bucket policy that requires Multi-Factor Authentication for requests to access the S3 bucket objects
- D. Enable server side encryption on the S3 bucket for data protection against the non-paying website visitors

Answer: B

NEW QUESTION 224

After launching an instance that you intend to serve as a NAT (Network Address Translation) device in a public subnet you modify your route tables to have the NAT device be the target of internet bound traffic of your private subnet. When you try and make an outbound connection to the Internet from an instance in the private subnet, you are not successful.

Which of the following steps could resolve the issue?

- A. Attaching a second Elastic Network interface (ENI) to the NAT instance, and placing it in the private subnet
- B. Attaching a second Elastic Network Interface (ENI) to the instance in the private subnet, and placing it in the public subnet
- C. Disabling the Source/Destination Check attribute on the NAT instance
- D. Attaching an Elastic IP address to the instance in the private subnet

Answer: C

NEW QUESTION 227

What happens, by default, when one of the resources in a CloudFormation stack cannot be created?

- A. Previously-created resources are kept but the stack creation terminates.
- B. Previously-created resources are deleted and the stack creation terminates.
- C. The stack creation continues, and the final results indicate which steps failed.
- D. CloudFormation templates are parsed in advance so stack creation is guaranteed to succeed

Answer: B

NEW QUESTION 228

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AWS-Certified-Developer-Associate Exam with Our Prep Materials Via below:

<https://www.certleader.com/AWS-Certified-Developer-Associate-dumps.html>