

➤ **Vendor: Amazon**

➤ **Exam Code: AWS Certified Developer - Associate**

➤ **Exam Name: Amazon AWS Certified Developer - Associate**

➤ **Question 1 – Question 30**

[Visit PassLeader and Download Full Version AWS Certified Developer - Associate Exam Dumps](#)

QUESTION 1

A user plans to use RDS as a managed DB platform. Which of the below mentioned features is not supported by RDS?

- A. Automated backup
- B. Automated scaling to manage a higher load
- C. Automated failure detection and recovery
- D. Automated software patching

Answer: B

Explanation:

AWS RDS provides a managed DB platform, which offers features, such as automated backup, patch management, automated failure detection and recovery. The scaling is not automated and the user needs to plan it with a few clicks.
<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html>

QUESTION 2

A user has not enabled versioning on an S3 bucket. What will be the version ID of the object inside that bucket?

- A. 0
- B. There will be no version attached
- C. Null
- D. Blank

Answer: C

Explanation:

S3 objects stored in the bucket before the user has set the versioning state have a version ID of null. When the user enables versioning, the objects in the bucket do not change and their ID remains null.
<http://docs.aws.amazon.com/AmazonS3/latest/dev/AddingObjectstoVersionSuspendedBuckets.html>

QUESTION 3

A user has created a queue named "myqueue" with SQS. There are four messages published to queue which are not received by the consumer yet. If the user tries to delete the queue, what will happen?

- A. A user can never delete a queue manually. AWS deletes it after 30 days of inactivity on queue
- B. It will initiate the delete but wait for four days before deleting until all messages are deleted automatically
- C. It will ask user to delete the messages first
- D. It will delete the queue

Answer: D

Explanation:

SQS allows the user to move data between distributed components of applications so they can perform different tasks without losing messages or requiring each component to be always available. The user can delete a queue at any time,

whether it is empty or not. It is important to note that queues retain messages for a set period of time. By default, a queue retains messages for four days.

<http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/SQSConcepts.html>

QUESTION 4

What happens if your application performs more reads or writes than your provisioned capacity?

- A. Nothing.
- B. requests above your provisioned capacity will be performed but you will receive 400 error codes.
- C. requests above your provisioned capacity will be performed but you will receive 200 error codes.
- D. requests above your provisioned capacity will be throttled and you will receive 400 error codes.

Answer: D

Explanation:

Speaking about DynamoDB, if your application performs more reads/second or writes/second than your table's provisioned throughput capacity allows, requests above your provisioned capacity will be throttled and you will receive 400 error codes.

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ProvisionedThroughputIntro.html>

QUESTION 5

In relation to Amazon SQS, how can you ensure that messages are delivered in order?

- A. Increase the size of your queue
- B. Send them with a timestamp
- C. Give each message a unique id
- D. AWS cannot guarantee that you will receive messages in the exact order you sent them

Answer: D

Explanation:

Amazon SQS makes a best effort to preserve order in messages, but due to the distributed nature of the queue, AWS cannot guarantee that you will receive messages in the exact order you sent them. You typically place sequencing information or timestamps in your messages so that you can reorder them upon receipt.

<https://aws.amazon.com/items/1343?externalID=1343>

QUESTION 6

An organization has launched two applications: one for blogging and one for ECM on the same AWS Linux EC2 instance running in the AWS VPC. The organization has attached two private IPs (primary and secondary) to the above mentioned instance. The organization wants the instance OS to recognize the secondary IP address. How can the organization configure this?

- A. Use the ec2-net-utility package which updates routing tables, uses DHCP to refresh the secondary IP and adds the network interface
- B. Use the ec2-net-utils package which will configure an additional network interface and update the routing table
- C. Use the ec2-ip-update package which can configure the network interface as well as update the secondary IP with DHCP
- D. Use the ec2-ip-utility package which can update the routing tables as well as refresh the secondary IP using DHCP

Answer: B

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. With VPC the user can specify multiple private IP addresses for his instances. The number of network interfaces and private IP addresses that a user can specify for an instance depends on the instance type. This scenario helps when the user wants to host multiple websites on a single EC2 instance. After the user has assigned a secondary private IP address to his instance, he needs to configure the operating system on that instance to recognize the secondary private IP address. For AWS Linux, the ec2-net-utils package can take care of this step. It configures additional network interfaces that the user can attach while the instance is running, refreshes secondary IP addresses during DHCP lease renewal, and updates the related routing rules.

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

QUESTION 7

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.
<http://aws.amazon.com/dynamodb/>

QUESTION 8

In relation to Amazon SQS, how many queues and messages can you have per queue for each user?

- A. Unlimited
- B. 10
- C. 256
- D. 500

Answer: A

Explanation:

Amazon SQS supports an unlimited number of queues and unlimited number of messages per queue for each user. Please be aware that Amazon SQS automatically deletes messages that have been in the queue for more than 4 days.
<https://aws.amazon.com/items/1343?externalID=1343>

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 purposes.

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
<http://aws.amazon.com/vpc/faqs/>

QUESTION 10

Regarding Amazon SQS, are there restrictions on the names of Amazon SQS queues?

- A. No
- B. Yes. Queue names must be unique within an AWS account and you cannot use hyphens (-) and underscores (_)
- C. Yes. Queue names are limited to 80 characters and queue names must be unique within an AWS account
- D. Yes. Queue names are limited to 80 characters but queue names do not need to be unique within an AWS account

Answer: C

Explanation:

Queue names are limited to 80 characters. Alphanumeric characters plus hyphens (-) and underscores (_) are allowed. Queue names must be unique within an AWS account. After you delete a queue, you can reuse the queue name.
<https://aws.amazon.com/sqs/faqs/>

QUESTION 11

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

- C. Launch a VPC instance with two network interfaces. Assign a separate security group and elastic IP to them.
- D. Launch a VPC with ELB such that it redirects requests to separate VPC instances of the public subnet.

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.

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

QUESTION 12

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.

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

QUESTION 13

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.

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

QUESTION 14

Regarding Amazon SWF, the coordination logic in a workflow is contained in a software program called a _____.

- A. Handler
- B. Decider
- C. Coordinator
- D. Worker

Answer: B

Explanation:

In Amazon SWF, the coordination logic in a workflow is contained in a software program called a decider. The decider schedules activity tasks, provides input data to the activity workers, processes events that arrive while the workflow is in progress, and ultimately ends (or closes) the workflow when the objective has been completed.

<http://docs.aws.amazon.com/amazonswf/latest/developerguide/swf-dg-intro-to-swf.html>

QUESTION 15

A user has attached one RDS security group with 5 RDS instances. The user has changed the ingress rule for the security group. What will be the initial status of the ingress rule?

- A. Approving
- B. Implementing
- C. Authorizing
- D. It is not possible to assign a single group to multiple DB instances

Answer: C

Explanation:

When the user makes any changes to the RDS security group the rule status will be authorizing for some time until the changes are applied to all instances that the group is connected with. Once the changes are propagated the rule status will change to authorized.

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithSecurityGroups.html

QUESTION 16

A user has attached an EBS volume to a running Linux instance as a "/dev/sdf" device. The user is unable to see the attached device when he runs the command "df -h". What is the possible reason for this?

- A. The volume is not in the same AZ of the instance
- B. The volume is not formatted
- C. The volume is not attached as a root device
- D. The volume is not mounted

Answer: D

Explanation:

When a user creates an EBS volume and attaches it as a device, it is required to mount the device. If the device/volume is not mounted it will not be available in the listing.

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

QUESTION 17

An account owner has created an IAM user with the name passleader. The account owner wants to give EC2 access of only the US West region to that IAM user. How can the owner configure this?

- A. While creating a policy provide the region as a part of the resources
- B. Create an IAM user in the US West region and give access to EC2
- C. Create an IAM policy and define the region in the condition
- D. It is not possible to provide access based on the region

Answer: C

Explanation:

The IAM policy is never region specific. If the user wants to configure the region specific setting, he needs to provide conditions as part of the policy.

<http://awspolicygen.s3.amazonaws.com/policygen.html>

QUESTION 18

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.

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

QUESTION 19

In DynamoDB, the default table size is:

- A. 5 GB
- B. 1 GB
- C. 10 GB
- D. There is no table size

Answer: D

Explanation:

DynamoDB has seamless scalability with no table size limits and unlimited storage, so you shouldn't be worried about managing storage on the host or to provisioning more drive, as your data requirement changes.

<http://aws.amazon.com/dynamodb/>

QUESTION 20

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.

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

QUESTION 21

You want to have multiple versions of your application running at the same time, with all versions launched via AWS Elastic Beanstalk. Is this possible?

- A. No. However if you have 2 AWS accounts this can be done
- B. No. AWS Elastic Beanstalk is not designed to support multiple running environments
- C. Yes. AWS Elastic Beanstalk is designed to support a number of multiple running environments
- D. Yes. However AWS Elastic Beanstalk is designed to support only 2 multiple running environments

Answer: C

Explanation:

AWS Elastic Beanstalk is designed to support multiple running environments. As an example you could have one for integration testing, one for pre-production, and one for production, with each environment independently configured and running on its own separate AWS resources.

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

QUESTION 22

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.

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

QUESTION 23

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.

<https://aws.amazon.com/items/1343?externalID=1343>

QUESTION 24

A user has launched one EC2 instance in the US West region. The user wants to access the RDS instance launched in the US East region from that EC2 instance. How can the user configure the access for that EC2 instance?

- A. It is not possible to access RDS of the US East region from the US West region
- B. Open the security group of the US West region in the RDS security group's ingress rule
- C. Configure the IP range of the US West region instance as the ingress security rule of RDS
- D. Create an IAM role which has access to RDS and launch an instance in the US West region with it

Answer: C

Explanation:

The user cannot authorize an Amazon EC2 security group if it is in a different AWS Region than the RDS DB instance. The user can authorize an IP range or specify an Amazon EC2 security group in the same region that refers to an IP address in another region.

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithSecurityGroups.html

QUESTION 25

In regard to AWS CloudFormation, what is a stack?

- A. The set of AWS templates that are created and managed as a template
- B. The set of AWS resources that are created and managed as a template
- C. The set of AWS resources that are created and managed as a single unit
- D. The set of AWS templates that are created and managed as a single unit

Answer: C

Explanation:

A stack is the set of AWS resources that are created and managed as a single unit when AWS CloudFormation initiates a template.

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/concept-stack.html>

QUESTION 26

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.

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DataModel.html>

QUESTION 27

Which of the below mentioned options is not a best practice to securely manage the AWS access credentials?

- A. Enable MFA for privileged users
- B. Create individual IAM users

- C. Keep rotating your secure access credentials at regular intervals
- D. Create strong access key and secret access key and attach to the root account

Answer: D

Explanation:

It is a recommended approach to avoid using the access and secret access keys of the root account. Thus, do not download or delete it. Instead make the IAM user as powerful as the root account and use its credentials. The user cannot generate their own access and secret access keys as they are always generated by AWS.
<http://docs.aws.amazon.com/IAM/latest/UserGuide/IAMBestPractices.html>

QUESTION 28

You have been given a scope to deploy some AWS infrastructure for a large organisation. 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.
<http://aws.amazon.com/autoscaling/>

QUESTION 29

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.
<http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/Welcome.html>

QUESTION 30

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 IP.

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.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html>

[Visit PassLeader and Download Full Version AWS Certified Developer - Associate Exam Dumps](#)