

Exam Questions DVA-C01

AWS Certified Developer Associate Exam

<https://www.2passeasy.com/dumps/DVA-C01/>



NEW QUESTION 1

In DynamoDB, if you create a table and request 10 units of write capacity and 200 units of read capacity of provisioned throughput, how much would you be charged in US East (Northern Virginia) Region?

- A. \$0.05 per hour
- B. \$0.10 per hour
- C. \$0.03 per hour
- D. \$0.15 per hour

Answer: A

Explanation:

To understand pricing in DynamoDB, consider the following example. If you create a table and request 10 units of write capacity and 200 units of read capacity of provisioned throughput, you would be charged:

$\$0.01 + (4 \times \$0.01) = \$0.05$ per hour

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

NEW QUESTION 2

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 3

A user has launched five instances with ELB. How can the user add the sixth EC2 instance to ELB?

- A. The user can add the sixth instance on the fly.
- B. The user must stop the ELB and add the sixth instance.
- C. The user can add the instance and change the ELB config file.
- D. The ELB can only have a maximum of five instance

Answer: A

Explanation:

Elastic Load Balancing automatically distributes incoming traffic across multiple EC2 instances. You create a load balancer and register instances with the load balancer in one or more Availability Zones. The load balancer serves as a single point of contact for clients. This enables you to increase the availability of your application. You can add and remove EC2 instances from your load balancer as your needs change, without disrupting the overall flow of information.

Reference: <http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/SvcIntro.htm>

NEW QUESTION 4

How long can you keep your Amazon SQS messages in Amazon SQS queues?

- A. From 120 secs up to 4 weeks
- B. From 10 secs up to 7 days
- C. From 60 secs up to 2 weeks
- D. From 30 secs up to 1 week

Answer: C

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 5

A user is planning to make a mobile game which can be played online or offline and will be hosted on EC2.

The user wants to ensure that if someone breaks the highest score or they achieve some milestone they can inform all their colleagues through email. Which of the below mentioned AWS services helps achieve this goal?

- A. AWS Simple Workflow Service.
- B. AWS Simple Queue Service.
- C. Amazon Cognito
- D. AWS Simple Email Service

Answer: D

Explanation:

Amazon Simple Email Service (Amazon SES) is a highly scalable and cost-effective email-sending service for businesses and developers. It integrates with other AWS services, making it easy to send emails from applications that are hosted on AWS.

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

NEW QUESTION 6

Which one of the following operations is NOT a DynamoDB operation?

- A. BatchWriteItem
- B. DescribeTable
- C. BatchGetItem
- D. BatchDeleteItem

Answer: D

Explanation:

In DynamoDB, DeleteItem deletes a single item in a table by primary key, but BatchDeleteItem doesn't exist.

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

NEW QUESTION 7

When a user is detaching an EBS volume from a running instance and attaching it to a new instance, which of the below mentioned options should be followed to avoid file system damage?

- A. Unmount the volume first
- B. Stop all the I/O of the volume before processing
- C. Take a snapshot of the volume before detaching
- D. Force Detach the volume to ensure that all the data stays intact

Answer: A

Explanation:

When a user is trying to detach an EBS volume, the user can either terminate the instance or explicitly remove the volume. It is a recommended practice to unmount the volume first to avoid any file system damage.

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

NEW QUESTION 8

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 9

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
- B. AWS deletes it after 30 days of inactivity on queue
- C. It will initiate the delete but wait for four days before deleting until all messages are deleted automatically.
- D. It will ask user to delete the messages first
- E. 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. Reference:

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

NEW QUESTION 10

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 code

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.

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

NEW QUESTION 10

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.

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

NEW QUESTION 12

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 14

Regarding Amazon SNS, to begin using Amazon SNS mobile push notifications, you first need that uses one of the supported push notification services: APNS, GCM, or ADM.

- A. an access policy for the mobile endpoints
- B. to active push notification service of Amazon SNS
- C. to know the type of mobile device operating system
- D. an app for the mobile endpoints

Answer: D

Explanation:

In Amazon SNS, to begin using Amazon SNS mobile push notifications, you first need an app for the mobile endpoints that uses one of the supported push notification services: APNS, GCM, or ADM. After you've registered and configured the app to use one of these services, you configure Amazon SNS to send push notifications to the mobile endpoints.

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

NEW QUESTION 19

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 22

Does AWS CloudFormation support Amazon EC2 tagging?

- A. It depends if the Amazon EC2 tagging has been defined in the template.

- B. No, it doesn't support Amazon EC2 tagging.
- C. No, CloudFormation doesn't support any tagging
- D. Yes, AWS CloudFormation supports Amazon EC2 tagging

Answer: D

Explanation:

In AWS CloudFormation, Amazon EC2 resources that support the tagging feature can also be tagged in an AWS template. The tag values can refer to template parameters, other resource names, resource attribute values (e.g. addresses), or values computed by simple functions (e.g., a concatenated list of strings).
Reference: <http://aws.amazon.com/cloudformation/faqs/>

NEW QUESTION 27

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

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 31

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 34

Can you configure an RDS Read Replica using CloudFormation templates?

- A. Yes, provided that you have root access.
- B. Yes, when you create a new CloudFormation template
- C. Yes, but not for all Regions.
- D. No, you can add the ReadReplica only when the resource is made available by CloudFormation

Answer: B

Explanation:

AWS CloudFormation gives developers and systems administrators an easy way to create and manage collections of AWS resources. You can now set Read Replicas for your databases with RDS when you create a new CloudFormation template. You can start using it with the sample template of CloudFormation.

Reference:

https://s3.amazonaws.com/cloudformation-templates-us-east-1/RDS_MySQL_With_Read_Replica.templ

NEW QUESTION 37

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. N
- B. However if you have 2 AWS accounts this can be done
- C. N
- D. AWS Elastic Beanstalk is not designed to support multiple running environments
- E. Ye
- F. AWS Elastic Beanstalk is designed to support a number of multiple running environments
- G. Ye
- H. 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.

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

NEW QUESTION 39

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 41

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 44

A user had defined an IAM policy similar to the one given below on a bucket:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Principal": {
      "AWS": "arn:aws:iam::12112112:user/test"
    }
  }],
  "Action": [ "s3:GetBucketLocation", "s3:ListBucket", "s3:GetObject" ],
  "Resource": [ "arn:aws:s3:::examkiller" ]
}
```

What will this do?

- A. It will result in an error saying invalid policy statement
- B. It will create an IAM policy for the user test
- C. Allows the user test of the AWS account ID 12112112 to perform GetBucketLocation, ListBucket and GetObject on the bucket examkiller
- D. It will allow all the IAM users of the account ID 12112112 to perform GetBucketLocation, ListBucket and GetObject on bucket examkiller

Answer: C

Explanation:

The IAM policy allows to test a user in the account 12112112 to perform: s3:GetBucketLocation

s3:ListBucket s3:GetObject

Amazon S3 permissions on the examkiller bucket.

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

NEW QUESTION 48

What does Amazon SQS provide?

- A. An asynchronous message queue service.
- B. A Simple Query Sewer, managed directly by Amazon Web Services.
- C. None of these.
- D. A synchronous message queue service.

Answer: A

Explanation:

Amazon SQS stands for Simple Queue Services, and provides a cost-effective way to decouple the components of your application through an asynchronous message queue service

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

NEW QUESTION 53

Bob is an IAM user who has access to the EC2 services. Admin is an IAM user who has access to all the AWS services including IAM. Can Bob change his password?

- A. No, the IAM user can never change the password
- B. Yes, provided Admin has given Bob access to change his password
- C. Yes, only from AWS CLI
- D. Yes, only from the AWS console

Answer: B

Explanation:

The IAM users by default cannot change their password. The root owner or IAM administrator needs to set the policy in the password policy page, which should allow the user to change their password. Once it is enabled, the IAM user can always change their passwords from the AWS console or CLI.

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

NEW QUESTION 55

Which of the below mentioned commands allows the user to share the AMI with his peers using the AWS EC2 CLI?

- A. ec2-share-image-public
- B. ec2-share-image-account
- C. ec2-share-image
- D. ec2-modify-image-attribute

Answer: D

Explanation:

A user can share an AMI with another user / peer using the command: ec2-modify-image-attribute <AMI-ID> -| -a <AWS Account ID>

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/sharingamis-explicit.html>

NEW QUESTION 57

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 like: <https://examkiller.signin.aws.amazon.com/console/>. How can this be configured?

- A. The user needs to use Route 53 to map the examkiller domain and IAM URL
- B. Create an IAM AWS account alias with the name examkiller
- C. It is not possible to have a personalized IAM login URL
- D. Create an IAM hosted zone Identity for the domain examkiller

Answer: B

Explanation:

If a user wants the URL of the AWS IAM sign-in page to have a company name instead of the AWS account ID, he can create an alias for his AWS account ID.

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

NEW QUESTION 58

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 60

An orgAMzation is having an application which can start and stop an EC2 instance as per schedule. The orgAMzation needs the MAC address of the instance to be registered with its software. The instance is launched in EC2-CLASSIC. How can the orgAMzation update the MAC registration every time an instance is booted?

- A. The instance MAC address never change
- B. Thus, it is not required to register the MAC address every time.
- C. The orgAMzation should write a boot strapping script which will get the MAC address from the instance metadata and use that script to register with the application.
- D. AWS never provides a MAC address to an instance; instead the instance ID is used for identifying the instance for any software registration.
- E. The orgAMzation should provide a MAC address as a part of the user dat
- F. Thus, whenever the instance is booted the script assigns the fixed MAC address to that instance.

Answer: B

Explanation:

AWS provides an on demand, scalable infrastructure. AWS EC2 allows the user to launch On-Demand instances. AWS does not provide a fixed MAC address to the instances launched in EC2-CLASSIC. If the instance is launched as a part of EC2-VPC, it can have an ENI which can have a fixed MAC. However, with EC2-CLASSIC, every time the instance is started or stopped it will have a new MAC address.

To get this MAC, the orgAMzation can run a script on boot which can fetch the instance metadata and get the MAC address from that instance metadata. Once the MAC is received, the orgAMzation can register that MAC with the software.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AESDg-chapter-instancedata.html>

NEW QUESTION 65

A user has created an EBS instance in the US-East-1a AZ. The user has a volume of 30 GB in the US-East-1 b zone. How can the user attach the volume to an instance?

- A. Since both the volume and the instance are in the same region, the user can attach the volume
- B. Use the volume migrate function to move the volume from one AZ to another and attach to the instance
- C. Take a snapshot of the volum
- D. Create a new volume in the USEast-1a and attach that to the instance
- E. Use the volume replicate function to create a new volume in the US-East-1a and attach that to the volume

Answer: C

Explanation:

If an EBS volume is not in the same AZ of an EC2 instance, it cannot be attached to the instance. The only option is to take a snapshot of the volume and create a new volume in the instance's AZ. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSSnapshots.html>

NEW QUESTION 68

You cannot access your AWS console, so you revert to using the CLI that you are not familiar with. Which of the following commands is not a valid CLI command for EC2 instances?

- A. ec2-allocate-address
- B. ec2-attach-internet-gateway
- C. ec2-associate-route-table
- D. ec2-allocate-interface

Answer: D

Explanation:

You can use the CLI tools to manage your Amazon EC2 resources (such as instances, security groups, and volumes) and your Amazon VPC resources (such as VPCs, subnets, route tables, and Internet gateways). Before you can start using the tools, you must download and configure them.

The following are valid CLI commands for EC2 instances: ec2-accept-vpc-peering-connection

ec2-allocate-address

ec2-assign-private-ip-addresses ec2-associate-address

ec2-associate-dhcp-options ec2-associate-route-table

ec2-attach-internet-gateway

ec2-attach-network-interface (not ec2-allocate-interface) Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/CommandLineReference/command-reference.html>

NEW QUESTION 72

In AWS Elastic Beanstalk, if the application returns any response other than 200 ,OK or there is no response within the configured InactMtyTimeout period, .

- A. SQS once again makes the message visible in the queue and available for another attempt at processing
- B. SQS waits for another timeout
- C. SQS run DeleteMessagecall and deletes the message from the queue
- D. SQS sends a message to the application with the IVlessageID and pending status

Answer: A

Explanation:

In AWS Elastic Beanstalk, if the application returns any response other than 200, OK or there is no response within the configured InactMtyTimeout period, SQS once again makes the message visible in the queue and available for another attempt at processing.

Reference:

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.html#worker-environment>

NEW QUESTION 75

Which header received at the EC2 instance identifies the port used by the client while requesting ELB?

- A. X-Forwarded-Proto
- B. X-Requested-Proto
- C. X-Forwarded-Port
- D. X-Requested-Port

Answer: C

Explanation:

The X-Forwarded-Port request header helps the user identify the port used by the client while sending a request to ELB.

Reference: <http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/TerminologyandKeyConcepts.html>

NEW QUESTION 79

A user has an S3 object in the US Standard region with the content "color=red". The user updates the object with the content as "color=white". If the user tries to read the value 1 minute after it was uploaded, what will S3 return?

- A. It will return "color=white"
- B. It will return "color=red"
- C. It will return an error saying that the object was not found
- D. It may return either "color=red" or "color=white" i.
- E. any of the value

Answer: D

Explanation:

AWS S3 follows the eventual consistent model in the US Standard Region. Once the object is updated it

may return the new value or the old value based on whether all the content is replicated across multiple servers until it becomes consistent (eventual).

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

NEW QUESTION 83

A user wants to access RDS from an EC2 instance using IP addresses. Both RDS and EC2 are in the same region, but different AZs. Which of the below mentioned options help configure that the instance is accessed faster?

- A. Configure the Private IP of the Instance in RDS security group
- B. Security group of EC2 allowed in the RDS security group
- C. Configuring the elastic IP of the instance in RDS security group
- D. Configure the Public IP of the instance in RDS security group

Answer: A

Explanation:

If the user is going to specify an IP range in RDS security group, AWS recommends using the private IP address of the Amazon EC2 instance. This provides a more direct network route from the Amazon EC2 instance to the Amazon RDS DB instance, and does not incur network charges for the data sent outside of the Amazon network.

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

NEW QUESTION 85

A user is creating a snapshot of an EBS volume. Which of the below statements is incorrect in relation to the creation of an EBS snapshot?

- A. Its incremental
- B. It can be used to launch a new instance
- C. It is stored in the same AZ as the volume
- D. It is a point in time backup of the EBS volume

Answer: C

Explanation:

The EBS snapshots are a point in time backup of the EBS volume. It is an incremental snapshot, but is always specific to the region and never specific to a single AZ.

Hence the statement "It is stored in the same AZ as the volume" is incorrect.

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

NEW QUESTION 90

Which of the following groups is AWS Elastic Beanstalk best suited for?

- A. Those who want to deploy and manage their applications within minutes in the AWS cloud
- B. Those who want to privately store and manage Git repositories in the AWS cloud.
- C. Those who want to automate the deployment of applications to instances and to update the applications as required
- D. Those who want to model, visualize, and automate the steps required to release software

Answer: A

Explanation:

AWS Elastic Beanstalk is best suited for those groups who want to deploy and manage their applications within minutes in the AWS cloud. As a bonus, you don't even need experience with cloud computing to get started.

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

NEW QUESTION 94

You are using Amazon SQS and are getting a "Queue Deleted Recently" error. What is wrong?

- A. The message is too big
- B. You have incorrect permissions
- C. Another user has deleted the queue
- D. If you delete a queue, you need to wait for at least 60 seconds before creating a queue with the same name

Answer: D

Explanation:

If you delete a queue, you need to wait for at least 60 seconds before creating a queue with the same name. Please note that when you delete a queue, the deletion process takes up to 60 seconds. Requests you send to a recently deleted queue might succeed during the 60-second period. For example, a `SendMessage` request might succeed, but after 60 seconds the queue and that message you sent no longer exists.

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

NEW QUESTION 96

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 97

A user is launching an AWS RDS with MySQL. Which of the below mentioned options allows the user to configure the INNODB engine parameters?

- A. Options group
- B. Engine parameters
- C. Parameter groups
- D. DB parameters

Answer: C

Explanation:

With regard to RDS, the user can manage the configuration of a DB engine by using a DB parameter group. A DB parameter group contains engine configuration values that can be applied to one or more DB instances of the same instance type.

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

NEW QUESTION 102

A user is setting up an Elastic Load Balancer(ELB). Which of the below parameters should the user consider so as the instance gets registered with the ELB?

- A. ELB DNS
- B. IP address
- C. Security group
- D. ELB IP

Answer: B

Explanation:

The EC2 instances are registered with the load balancer using the IP addresses associated with the instances. When an instance is stopped and then started, the IP address associated with the instance changes. This prevents the load balancer from routing traffic to the restarted instance. When the user stops and then starts registered EC2 instances, it is recommended that to de-register the stopped instance from load balancer, and then register the restarted instance. Failure to do so may prevent the load balancer from performing health checks and routing the traffic to the restarted instance.

NEW QUESTION 107

The user has configured AutoScaling based on the dynamic policy. Which of the following is not the right command to specify a change in capacity as a part of the policy?

- A. "adjustment=-50" (type is PercentChangeInCapacity)
- B. "adjustment=3" (type is ExactCapacity)
- C. "adjustment=-1" (type is ChangeInCapacity)
- D. "adjustment=-8" (type is ExactCapacity)

Answer: D

Explanation:

The user can configure the AutoScaling group to automatically scale up and then scale down based on the various specified CloudWatch monitoring conditions. The user needs to provide the adjustment value and the adjustment type. A positive adjustment value increases the current capacity and a negative adjustment value decreases the current capacity. The user can express the change to the current size as an absolute number, an increment or as a percentage of the current group size.

In this option specifying the exact capacity with the adjustment value = -8 will not work as when type is exact capacity the adjustment value cannot be negative.

Reference:

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

NEW QUESTION 112

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 instance
- 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 unhealthy
- 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 115

An organization has created an application which is hosted on the AWS EC2 instance. The application stores images to S3 when the end user uploads to it. The organization does not want to store the AWS secure credentials required to access the S3 inside the instance. Which of the below mentioned options is a possible solution to avoid any security threat?

- A. Use the IAM role and assign it to the instance.
- B. Since the application is hosted on EC2, it does not need credentials to access S3.
- C. Use the X.509 certificates instead of the access and the secret access keys.
- D. Use the IAM based single sign between the AWS resources and the organization application

Answer: A

Explanation:

The AWS IAM role uses temporary security credentials to access AWS services. Once the role is assigned to an instance, it will not need any security credentials to be stored on the instance. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

NEW QUESTION 116

AWS Elastic Beanstalk stores your application files and optionally server log files in .

- A. Amazon Storage Gateway
- B. Amazon Glacier
- C. Amazon EC2
- D. Amazon S3

Answer: D

Explanation:

AWS Elastic Beanstalk stores your application files and optionally server log files in Amazon S3. If you are using the AWS Management Console, Git, the AWS Toolkit for Visual Studio, or AWS Toolkit for Eclipse, an Amazon S3 bucket will be created in your account for you and the files you upload will be automatically copied from your local client to Amazon S3. Optionally, you may configure Elastic Beanstalk to copy your server log files every hour to Amazon S3. You do this by editing the environment configuration settings.

Reference: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/AWSHowTo.html>

NEW QUESTION 118

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 122

ExamKiller has three AWS accounts. They have created separate IAM users within each account. ExamKiller wants a single IAM console URL such as <https://examkiller.signin.aws.amazon.com/console/> for all account users. How can this be achieved?

- A. Merge all the accounts with consolidated billing
- B. Create the same account alias with each account ID
- C. It is not possible to have the same IAM account login URL for separate AWS accounts
- D. Create the S3 bucket with an alias name and use the redirect rule to forward requests to various accounts

Answer: C

Explanation:

If a user wants the URL of the AWS IAM sign-in page to have a company name instead of the AWS account ID, he can create an alias for his AWS account ID. The alias should be unique.

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

NEW QUESTION 124

An EC2 instance has one additional EBS volume attached to it. How can a user attach the same volume to another running instance in the same AZ?

- A. Terminate the first instance and only then attach to the new instance
- B. Attach the volume as read only to the second instance
- C. Detach the volume first and attach to new instance
- D. No need to detach
- E. Just select the volume and attach it to the new instance, it will take care of mapping internally

Answer: C

Explanation:

If an EBS volume is attached to a running EC2 instance, the user needs to detach the volume from the original instance and then attach it to a new running instance. The user doesn't need to stop / terminate the original instance.

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

NEW QUESTION 129

When working with AWS CloudFormation Templates what is the maximum number of stacks that you can create?

- A. 500
- B. 50
- C. 20
- D. 10

Answer: C

Explanation:

CloudFormation Limits

Maximum number of AWS CloudFormation stacks that you can create is 20 stacks. Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cloudformation-limits.html>

NEW QUESTION 131

A bucket owner has allowed another account's IAM users to upload or access objects in his bucket. The IAM user of Account A is trying to access an object created by the IAM user of account B. What will happen in this scenario?

- A. The bucket policy may not be created as S3 will give error due to conflict of Access Rights
- B. It is not possible to give permission to multiple IAM users
- C. AWS S3 will verify proper rights given by the owner of Account A, the bucket owner as well as by the IAM user B to the object
- D. It is not possible that the IAM user of one account accesses objects of the other IAM user

Answer: C

Explanation:

If a IAM user is trying to perform some action on an object belonging to another AWS user's bucket, S3 will verify whether the owner of the IAM user has given sufficient permission to him. It also verifies the policy for the bucket as well as the policy defined by the object owner.

Reference:

<http://docs.aws.amazon.com/AmazonS3/latest/dev/access-control-auth-workflow-object-operation.html>

NEW QUESTION 133

How can you secure data at rest on an EBS volume?

- A. Attach the volume to an instance using EC2's SSL interface.
- B. Write the data randomly instead of sequentially.
- C. Use an encrypted file system on top of the EBS volume.
- D. Encrypt the volume using the S3 server-side encryption service.
- E. Create an IAM policy that restricts read and write access to the volume

Answer: C

NEW QUESTION 137

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 141

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 144

What is one key difference between an Amazon EBS-backed and an instance-store backed instance?

- A. Virtual Private Cloud requires EBS backed instances
- B. Amazon EBS-backed instances can be stopped and restarted
- C. Auto scaling requires using Amazon EBS-backed instances.
- D. Instance-store backed instances can be stopped and restarte

Answer: B

NEW QUESTION 148

You have written an application that uses the Elastic Load Balancing service to spread traffic to several web servers Your users complain that they are sometimes forced to login again in the middle of using your application, after they have already toggged in. This is not behavior you have designed. What is a possible solution to prevent this happening?

- A. Use instance memory to save session state.
- B. Use instance storage to save session state.
- C. Use EBS to save session state
- D. Use ElastiCache to save session state.
- E. Use Glacier to save session slat

Answer: D

NEW QUESTION 150

EC2 instances are launched from Amazon Machine images (AMIS). A given public AMI can:

- A. be used to launch EC2 Instances in any AWS region.
- B. only be used to launch EC2 instances in the same country as the AMI is stored.
- C. only be used to launch EC2 instances in the same AWS region as the AMI is stored.
- D. only be used to launch EC2 instances in the same AWS availability zone as the AMI is stored

Answer: C

NEW QUESTION 155

Which of the following platforms are supported by Elastic Beanstalk? Choose 2 answers

- A. Apache Tomcat
- B. .NET
- C. IBM Websphere
- D. Oracle JBoss
- E. Jetty

Answer: AB

NEW QUESTION 157

Which EC2 API call would you use to retrieve a list of Amazon Machine Images (AMIs)?

- A. DescnbelInstances
- B. DescribeAMIs
- C. Describelmages
- D. GetAMIs
- E. You cannot retrieve a list of AMIs as there are over 10,000 AMIs

Answer: E

NEW QUESTION 161

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 162

Company B provides an online image recognition service and utilizes SQS to decouple system components for scalability. The SQS consumers poll the imaging queue as often as possible to keep end-to-end throughput as high as possible. However, Company B is realizing that polling in tight loops is burning CPU cycles and increasing costs with empty responses.

How can Company B reduce the number of empty responses?

- A. Set the imaging queue visibility Timeout attribute to 20 seconds
- B. Set the Imaging queue ReceiveMessageWaitTimeSeconds attribute to 20 seconds
- C. Set the imaging queue MessageRetentionPeriod attribute to 20 seconds
- D. Set the DelaySeconds parameter of a message to 20 seconds

Answer: B

NEW QUESTION 165

How can software determine the public and private IP addresses of the Amazon EC2 instance that it is running on?

- A. Query the appropriate Amazon CloudWatch metric.
- B. Use ipconfig or ifconfig command.
- C. Query the local instance userdata.
- D. Query the local instance metadata

Answer: D

NEW QUESTION 170

Which of the following items are required to allow an application deployed on an EC2 instance to write data to a DynamoDB table? Assume that no security keys are allowed to be stored on the EC2 instance. Choose 2 answers

- A. Create an IAM User that allows write access to the DynamoDB table.
- B. Add an IAM Role to a running EC2 instance.
- C. Add an IAM User to a running EC2 Instance.
- D. Launch an EC2 Instance with the IAM Role included in the launch configuration.
- E. Create an IAM Role that allows write access to the DynamoDB table.
- F. Launch an EC2 Instance with the IAM User included in the launch configuration.

Answer: DE

NEW QUESTION 175

How is provisioned throughput affected by the chosen consistency model when reading data from a DynamoDB table?

- A. Strongly consistent reads use the same amount of throughput as eventually consistent reads
- B. Strongly consistent reads use more throughput than eventually consistent reads.
- C. Strongly consistent reads use less throughput than eventually consistent reads
- D. Strongly consistent reads use variable throughput depending on read actMty

Answer: B

NEW QUESTION 180

A corporate web application is deployed within an Amazon VPC, and is connected to the corporate data center via IPsec VPN. The application must authenticate against the on-premise LDAP server. Once authenticated, logged-in users can only access an S3 keyspace specific to the user. Which two approaches can satisfy the objectives? Choose 2 answers

- A. The application authenticates against LDA
- B. The application then calls the IAM Security Service to login to IAM using the LDAP credential
- C. The application can use the IAM temporary credentials to access the appropriate S3 bucket.
- D. The application authenticates against LDAP, and retrieves the name of an IAM role associated with the use
- E. The application then calls the IAM Security Token Service to assume that IAM Rol
- F. The application can use the temporary credentials to access the appropriate S3 bucket.
- G. The application authenticates against IAM Security Token Service using the LDAP credential
- H. The application uses those temporary AWS security credentials to access the appropriate S3 bucket.
- I. Develop an identity broker which authenticates against LDAP, and then calls IAM Security Token Service to get IAM federated user credential
- J. The application calls the identity broker to get IAM federated user credentials with access to the appropriate S3 bucket.
- K. Develop an identity broker which authenticates against IAM Security Token Service to assume an IAM Role to get temporary AWS security credential
- L. The application calls the identity broker to get AWS temporary security credentials with access to the appropriate S3 bucket.

Answer: BD

NEW QUESTION 185

You attempt to store an object in the US-STANDARD region in Amazon S3, and receive a confirmation that it has been successfully stored. You then immediately make another API call and attempt to read this object. S3 tells you that the object does not exist. What could explain this behavior?

- A. US-STANDARD uses eventual consistency and it can take time for an object to be readable in a bucket
- B. Objects in Amazon S3 do not become visible until they are replicated to a second region.
- C. US-STANDARD imposes a 1 second delay before new objects are readable.
- D. You exceeded the bucket object limit, and once this limit is raised the object will be visible.

Answer: A

NEW QUESTION 190

You are writing to a DynamoDB table and receive the following exception: "ProvisionedThroughputExceededException". though according to your Cloudwatch metrics for the table, you are not exceeding your provisioned throughput. What could be an explanation for this?

- A. You haven't provisioned enough DynamoDB storage instances
- B. You're exceeding your capacity on a particular Range Key
- C. You're exceeding your capacity on a particular Hash Key
- D. You're exceeding your capacity on a particular Sort Key
- E. You haven't configured DynamoDB Auto Scaling triggers

Answer: C

NEW QUESTION 193

If an application is storing hourly log files from thousands of instances from a high traffic web site, which naming scheme would give optimal performance on S3?

- A. Sequential
- B. instanceID_log-HH-DD-NIM-YYYY
- C. instanceID_log-YYYY-NIM-DD-HH
- D. HH-DD-NINI-YYYY-log_instanceID
- E. YYYY-MM-DD-HH-|log_instanceID

Answer: E

NEW QUESTION 196

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DVA-C01 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DVA-C01 Product From:

<https://www.2passeasy.com/dumps/DVA-C01/>

Money Back Guarantee

DVA-C01 Practice Exam Features:

- * DVA-C01 Questions and Answers Updated Frequently
- * DVA-C01 Practice Questions Verified by Expert Senior Certified Staff
- * DVA-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * DVA-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year