

Exam Questions DVA-C01

AWS Certified Developer Associate Exam

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



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

An organization has 500 employees. The organization wants to set up AWS access for each department. Which of the below mentioned options is a possible solution?

- A. Create IAM roles based on the permission and assign users to each role
- B. Create IAM users and provide individual permission to each
- C. Create IAM groups based on the permission and assign IAM users to the groups
- D. It is not possible to manage more than 100 IAM users with AWS

Answer: C

Explanation:

An IAM group is a collection of IAM users. Groups let the user specify permissions for a collection of users, which can make it easier to manage the permissions for those users.

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

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

In regard to DynamoDB, which of the following statements is correct?

- A. An Item should have at least two value sets, a primary key and another attribute.
- B. An Item can have more than one attributes.
- C. A primary key should be single-valued.
- D. An attribute can have one or several other attribute

Answer: B

Explanation:

In Amazon DynamoDB, a database is a collection of tables. A table is a collection of items and each item is a collection of attributes.

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

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

In Amazon EC2, which of the following is the type of monitoring data for Amazon EBS volumes that is available automatically in 5-minute periods at no charge?

- A. Primary
- B. Basic
- C. Initial
- D. Detailed

Answer: B

Explanation:

Basic is the type of monitoring data (for Amazon EBS volumes) which is available automatically in 5-minute periods at no charge called.

Reference:

<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/monitoring-volume-status.html>

NEW QUESTION 8

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 9

A user is planning to create a structured database in the cloud. Which of the below mentioned AWS offerings help the user achieve the goal?

- A. AWS DynamoDB
- B. AWS RDS
- C. AWS SimpleDB
- D. AWS RSD

Answer: B

Explanation:

AWS RDS is a managed database server offered by AWS, which makes it easy to set up, operate, and scale a relational database or structured data in cloud.

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

NEW QUESTION 10

A user has created a MySQL RDS instance with PIOPS. Which of the below mentioned statements will help user understand the advantage of PIOPS?

- A. The user can achieve additional dedicated capacity for the EBS I/O with an enhanced RDS option
- B. It uses optimized EBS volumes and optimized configuration stacks
- C. It provides a dedicated network bandwidth between EBS and RDS
- D. It uses a standard EBS volume with optimized configuration the stacks

Answer: B

Explanation:

RDS DB instance storage comes in two types: standard and provisioned IOPS. Standard storage is allocated on the Amazon EBS volumes and connected to the

user's DB instance. Provisioned IOPS uses optimized EBS volumes and an optimized configuration stack. It provides additional, dedicated capacity for the EBS I/O.

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

NEW QUESTION 10

A user is accessing an EC2 instance on the SSH port for IP 10.20.30.40. Which one is a secure way to configure that the instance can be accessed only from this IP?

- A. In the security group, open port 22 for IP 1020.30.40/0
- B. In the security group, open port 22 for IP 10.20.30.40/32
- C. In the security group, open port 22 for IP 10.20.30.40/24
- D. In the security group, open port 22 for IP 10.20.30.40

Answer: B

Explanation:

In AWS EC2, while configuring a security group, the user needs to specify the IP address in CIDR notation. The CIDR IP range 10.20.30.40/32 says it is for a single IP 10.20.30.40. If the user specifies the IP as 10.20.30.40 only, the security group will not accept and ask it in a CIRD format.

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

NEW QUESTION 15

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.

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

NEW QUESTION 16

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 21

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.

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

NEW QUESTION 23

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

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

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

NEW QUESTION 26

In Amazon SNS, to send push notifications to mobile devices using Amazon SNS and ADM, you need to obtain the following, except:

- A. Client secret
- B. Client ID
- C. Device token
- D. Registration ID

Answer: C

Explanation:

To send push notifications to mobile devices using Amazon SNS and ADM, you need to obtain the following: Registration ID and Client secret.

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

NEW QUESTION 31

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 34

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.amazonaws.com/AWSEC2/latest/UserGuide/block-device-mapping-concepts.html>

NEW QUESTION 38

Can a user get a notification of each instance start / terminate configured with Auto Scaling?

- A. Yes, always
- B. No
- C. Yes, if configured with the Auto Scaling group
- D. Yes, if configured with the Launch Config

Answer: C

Explanation:

The user can get notifications using SNS if he has configured the notifications while creating the Auto Scaling group.

Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/GettingStartedTutorial.html>

NEW QUESTION 40

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 45

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 46

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.

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

NEW QUESTION 50

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 55

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 58

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 receive 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 63

In AWS Elastic Beanstalk, you can update your deployed application even while it is part of a running environment. For a Java application, you can also use to update your deployed application.

- A. the AWS Toolkit for Eclipse
- B. the AWS Toolkit for Visual Studio
- C. the AWS Toolkit for JVM
- D. the AWS Toolkit for Netbeans

Answer: A

Explanation:

In AWS Elastic Beanstalk, you can update your deployed application, even while it is part of a running environment. For a Java application, you can also use the AWS Toolkit for Eclipse to update your deployed application.

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

NEW QUESTION 67

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.template

NEW QUESTION 69

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.

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

NEW QUESTION 70

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 72

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 73

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.

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

NEW QUESTION 76

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 77

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 78

A user has configured a bucket S3 to host a static website. What difference will there be when static website hosting is enabled?

- A. It will help the user identify this bucket as the website root to map with the domain
- B. It will create a new version of the bucket
- C. It will not make any difference, but will help the user to configure the error page
- D. It will provide the region specific website endpoint

Answer: D

Explanation:

To host a static website, the user needs to configure an Amazon S3 bucket for website hosting and then upload the website contents to the bucket. The website is then available at the region-specific website endpoint of the bucket.

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

NEW QUESTION 82

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 operations.

- A. None of the above
- B. Both
- C. Query
- D. Scan

Answer: C

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.

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

NEW QUESTION 86

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 91

In DynamoDB, could you use IAM to grant access to Amazon DynamoDB resources and API actions?

- A. Yes
- B. Depended to the type of access
- C. In DynamoDB there is no need to grant access
- D. No

Answer: A

Explanation:

Amazon DynamoDB integrates with AWS Identity and Access Management (IAM). You can use AWS IAM to grant access to Amazon DynamoDB resources and API actions. To do this, you first write an AWS IAM policy, which is a document that explicitly lists the permissions you want to grant. You then attach that policy to an AWS IAM user or role.

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

NEW QUESTION 92

A user is planning to host a mobile game on EC2 which sends notifications to active users on either high score or the addition of new features. The user should get this notification when he is online on his mobile device. Which of the below mentioned AWS services can help achieve this functionality?

- A. AWS Simple Notification Service.
- B. AWS Simple Queue Service.
- C. AWS Mobile Communication Service.
- D. AWS Simple Email Service

Answer: A

Explanation:

Amazon Simple Notification Service (Amazon SNS) is a fast, flexible, and fully managed push messaging service. Amazon SNS makes it simple and cost-effective to push to mobile devices, such as iPhone, iPad, Android, Kindle Fire, and internet connected smart devices, as well as pushing to other distributed services.

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

NEW QUESTION 93

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 95

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 98

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 101

A user has created a new EBS volume from an existing snapshot. The user mounts the volume on the instance to which it is attached. Which of the below mentioned options is a required step before the user can mount the volume?

- A. Run a cyclic check on the device for data consistency
- B. Create the file system of the volume
- C. Resize the volume as per the original snapshot size
- D. No step is required
- E. The user can directly mount the device

Answer: D

Explanation:

When a user is trying to mount a blank EBS volume, it is required that the user first creates a file system within the volume. If the volume is created from an existing snapshot then the user needs not to create a file system on the volume as it will wipe out the existing data.

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

NEW QUESTION 105

In regards to Amazon SQS how many times will you receive each message?

- A. At least twice
- B. Exactly once
- C. As many times as you want
- D. At least once

Answer: D

Explanation:

Amazon SQS is engineered to provide "at least once" delivery of all messages in its queues. Although most of the time, each message will be delivered to your application exactly once, you should design your system so that processing a message more than once does not create any errors or inconsistencies. Reference: <https://aws.amazon.com/sqs/faqs/>

NEW QUESTION 106

A user is enabling logging on a particular bucket. Which of the below mentioned options may be best suitable to allow access to the log bucket?

- A. Create an IAM policy and allow log access
- B. It is not possible to enable logging on the S3 bucket
- C. Create an IAM Role which has access to the log bucket
- D. Provide ACL for the logging group

Answer: D

Explanation:

The only recommended use case for the S3 bucket ACL is to grant the write permission to the Amazon S3 Log Delivery group to write access log objects to the user's bucket.

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

NEW QUESTION 107

A user is running a webserver on EC2. The user wants to receive the SMS when the EC2 instance utilization is above the threshold limit. Which AWS services should the user configure in this case?

- A. AWS CloudWatch + AWS SES.
- B. AWS CloudWatch + AWS SNS.
- C. AWS CloudWatch + AWS SQS.
- D. AWS EC2 + AWS Cloudwatc

Answer: B

Explanation:

Amazon SNS makes it simple and cost-effective to push to mobile devices, such as iPhone, iPad, Android, Kindle Fire, and internet connected smart devices, as well as pushing to other distributed services. In this case, the user can configure that Cloudwatch sends an alarm on when the threshold is crossed to SNS which will trigger an SMS.

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

NEW QUESTION 108

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 110

A user is trying to share a video file with all his friends. Which of the below mentioned AWS services will be cheapest and easy to use?

- A. AWS S3
- B. AWS EC2
- C. AWS RRS
- D. AWS Glacier

Answer: C

Explanation:

AWS RRS provides the same functionality as AWS S3, but at a cheaper rate. It is ideally suited for non mission critical applications. It provides less durability than S3, but is a cheaper option.

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

NEW QUESTION 114

A user has configured ELB. Which of the below mentioned protocols the user can configure for ELB health checks while setting up ELB?

- A. All of the options
- B. TCP
- C. HTTPS
- D. SSL

Answer: A

Explanation:

An ELB performs a health check on its instances to ensure that it diverts traffic only to healthy instances. The ELB can perform a health check on HTTP, HTTPS, TCP and SSL protocols.

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

NEW QUESTION 115

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 volume
- 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 119

An organization has 10000 employees. The organization wants to give restricted AWS access to each employee. How can the organization achieve this?

- A. Create an IAM user for each employee and make them a part of the group
- B. It is not recommended to support 10000 users with IAM
- C. Use STS and create the users' run time
- D. Use Identity federation with SSO

Answer: D

Explanation:

Identity federation enables users from an existing directory to access resources within your AWS account, making it easier to manage your users by maintaining their identities in a single place. In this case, the federated user is the only solution since AWS does not allow creating more than 5000 IAM users. Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html>

NEW QUESTION 121

Can a user associate and use his own DNS with ELB instead of the DNS provided by AWS ELB?

- A. Yes, by creating a CNAME with the existing domain name provider
- B. Yes, by configuring DNS in the AWS Console
- C. No
- D. Yes, only through Route 53 by mapping ELB and DNS

Answer: A

Explanation:

The AWS ELB allows mapping a custom domain name with ELB. The user can map ELB with DNS in two ways: 1) By creating CNAME with the existing domain name service provider or 2) By creating a record with Route 53.

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/using-domain-names-with-elb.html>

NEW QUESTION 122

can be used to bootstrap both the Chef Server and Chef Client software on your EC2 instances.

- A. AWS CloudFormation
- B. AWS Elastic Beanstalk
- C. AWS OpsWorks
- D. Amazon Glacier

Answer: A

Explanation:

AWS CloudFormation can be used to bootstrap both the Chef Server and Chef Client software on your EC2 instances.

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

NEW QUESTION 123

In relation to Amazon Simple Workflow Service (Amazon SWF), what is an "Activity Worker"?

- A. An individual task undertaken by a workflow
- B. The automation of a business process
- C. A piece of software that implements tasks
- D. All answers listed are correct

Answer: C

Explanation:

In relation to Amazon Simple Workflow Service (Amazon SWF), an activity worker is a program that receives activity tasks, performs them, and provides results back. Which translates to a piece of software that implements tasks.

Reference: <http://docs.aws.amazon.com/amazonswf/latest/developerguide/swf-dg-develop-act.html>

NEW QUESTION 125

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 130

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 132

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 134

Which Amazon service is not used by Elastic Beanstalk?

- A. Amazon S3
- B. Amazon ELB
- C. Auto scaling
- D. Amazon EMR

Answer: D

Explanation:

Elastic Beanstalk leverages AWS services such as Amazon Elastic Cloud Compute (Amazon EC2), Amazon Simple Storage Service (Amazon S3), Amazon Simple Notification Service (Amazon SNS), Elastic Load Balancing and Auto Scaling to deliver the same highly reliable, scalable, and cost-effective infrastructure that hundreds of thousands of businesses depend on today.

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

NEW QUESTION 138

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 141

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 142

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 147

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 152

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 155

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 156

What is the maximum number of tags that a user can assign to an EC2 instance?

- A. 50
- B. 10
- C. 5
- D. 25

Answer: B

Explanation:

To help manage EC2 instances as well as their usage in a better way, the user can tag the instances. The tags are metadata assigned by the user which consists of a key and a value. One resource can have a maximum of 10 tags.

Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html

NEW QUESTION 158

How do you configure SQS to support longer message retention?

- A. Set the `MessageRetentionPeriod` 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 `MessageRetentionPeriod` 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 `MessageRetentionPeriod` 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 162

An organization is hosting an application as part of the free usage tier. The organization wants to create IAM users for each of its 150 employees and they may access AWS as part of free usage tier. What will you advise the organization?

- A. The IAM is not available as a part of the free usage tier
- B. Create IAM roles and give access based on role since it will not cost the user
- C. Do not create more than 100 users as it will cost the organization.
- D. Create IAM users for each employee as it does not cost

Answer: D

Explanation:

IAM is a free service. You can create as many IAM users or groups as desired free of cost. Reference:

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

NEW QUESTION 166

A user has configured ELB with two instances running in separate AZs of the same region? Which of the below mentioned statements is true?

- A. Multiple AZ instances will provide HA with ELB
- B. Multiple AZ instances are not possible with a single ELB
- C. Multiple AZ instances will provide scalability with ELB
- D. The user can achieve both HA and scalability with ELB

Answer: A

Explanation:

If a user is running two instances in separate AZs, it will provide HA with ELB since ELB will automatically stop routing the traffic to unhealthy instances and send it to healthy instances only.

NEW QUESTION 167

A user has developed an application which is required to send the data to a NoSQL database. The user wants to decouple the data sending such that the application keeps processing and sending data but does not wait for an acknowledgement of DB. Which of the below mentioned applications helps in this scenario?

- A. AWS Simple Notification Service
- B. AWS Simple Workflow
- C. AWS Simple Query 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. In this case, the user can use AWS SQS to send messages which are received from an application and sent to DB. The application can continue processing data without waiting for any acknowledgement from DB. The user can use SQS to transmit any volume of data without losing messages or requiring other services to always be available.

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

NEW QUESTION 172

A user is configuring the HTTPS protocol on a front end ELB and the SSL protocol for the back-end listener in ELB. What will ELB do?

- A. It will allow you to create the configuration, but the instance will not pass the health check
- B. Receives requests on HTTPS and sends it to the back end instance on SSL
- C. It will not allow you to create this configuration
- D. It will allow you to create the configuration, but ELB will not work as expected

Answer: C

Explanation:

If a user is configuring HTTPS on the front end and TCP on the back end, ELB will not allow saving these listeners and will respond with the message.

"Load Balancer protocol is an application layer protocol, but instance protocol is not. Both the Load Balancer protocol and the instance protocol should be at the same layer. Please fix."

Reference:

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

NEW QUESTION 177

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 179

When you use the AWS Elastic Beanstalk console to deploy a new application .

- A. you'll need to upload each file separately
- B. you'll need to create each file and path
- C. you'll need to upload a source bundle
- D. you'll need to create each file

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.

Reference:

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.deployment.source.html>

NEW QUESTION 181

A user is planning to use the AWS RDS with MySQL. Which of the below mentioned services the user is not going to pay?

- A. Data transfer
- B. RDS Cloudwatch metrics
- C. Data storage
- D. I/O requests per month

Answer: B

Explanation:

RDS charges the user on a pay as you go basis. It charges the user based on the instance type, number of hours that the instance is running, data transfer, storage cost as well for the I/O requests. The monitoring is free of cost.

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

NEW QUESTION 184

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 185

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 Server 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 189

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 190

An organization has created multiple components of a single application for compartmentalization. Currently all the components are hosted on a single EC2 instance. Due to security reasons the organization wants to implement two separate VPCs for the separate modules although it is already using VPC. How can the organization achieve this with a single instance?

- A. Create a VPC instance which will have both the ACL and the security group attached to it and have separate rules for each IP address.
- B. Create a VPC instance which will have multiple network interfaces with multiple elastic IP addresses.
- C. You have to launch two instances each in a separate subnet and allow VPC peering for a single IP.
- D. Create a VPC instance which will have multiple subnets attached to it and each will have a separate IP address.

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. With each network interface the organization can assign an EIP. This scenario helps when the user wants to host multiple websites on a single EC2 instance by using multiple SSL certificates on a single server and associating each certificate with a specific EIP address. It also helps in scenarios 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 195

In regards to VPC, select the correct statement:

- A. You can associate multiple subnets with the same Route Table.
- B. You can associate multiple subnets with the same Route Table, but you can't associate a subnet with only one Route Table.
- C. You can't associate multiple subnets with the same Route Table.
- D. None of these

Answer: A

Explanation:

Every subnet in your VPC must be associated with exactly one Route Table. However, multiple subnets can be associated with the same Route Table.

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

NEW QUESTION 197

Which of the following device names is reserved for the root device for Linux instances of Amazon EC2?

- A. /dev/sda1
- B. /dev/sd[b-e]
- C. xvd[a-e]
- D. /dev/sd[f-p][1-6]

Answer: A

Explanation:

/dev/sda1 is the name of the device reserved for the root device for Linux instances. Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/device_naming.html

NEW QUESTION 200

A user has enabled automated backup for an RDS instance. What is the longest duration for which the user can retain the automated backup?

- A. 25 days
- B. 15 days
- C. 45 days
- D. 35 days

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. The maximum period can be 35 days.

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

NEW QUESTION 201

A user is uploading archives to Glacier. The user is trying to understand key Glacier resources. Which of the below mentioned options is not a Glacier resource?

- A. Notification configuration
- B. Archive ID
- C. Job
- D. Archive

Answer: B

Explanation:

AWS Glacier has four resources. Vault and Archives are core data model concepts. Job is required to initiate download of archive. The notification configuration is required to send user notification when archive is available for download.

Reference: <http://docs.aws.amazon.com/amazonglacier/latest/dev/amazon-glacier-data-model.html>

NEW QUESTION 203

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 205

Which of the following programming languages have an officially supported AWS SDK? Choose 2 answers

- A. Perl
- B. PHP
- C. Pascal
- D. Java
- E. SQL

Answer: BD

NEW QUESTION 207

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 209

Which of the following statements about SWF are true? Choose 3 answers

- A. SWF tasks are assigned once and never duplicated
- B. SWF requires an S3 bucket for workflow storage
- C. SWF workflow executions can last up to a year
- D. SWF triggers SNS notifications on task assignment
- E. SWF uses deciders and workers to complete tasks
- F. SWF requires at least 1 EC2 instance per domain

Answer: ACE

NEW QUESTION 211

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 requests are allowed

Answer: AC

NEW QUESTION 215

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-west-1.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-Disposition header to the request for web fonts 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 217

A startup's photo-sharing site is deployed in a VPC. An ELB distributes web traffic across two subnets. ELB session stickiness is configured to use the AWS-generated session cookie, with a session TTL of 5 minutes. The webserver Auto Scaling Group is configured as: min-size=4, max-size=4. The startup is preparing for a public launch, by running load-testing software installed on a single EC2 instance running in us-west-2a. After 60 minutes of load-testing, the webserver logs show:

| WEBSERVER LOGS | # of HTTP requests from load-tester | # of HTTP requests from private beta users |
|--------------------------------------|-------------------------------------|--|
| webserver #1 (subnet in us-west-2a): | 19,210 | 434 |
| webserver #2 (subnet in us-west-2a): | 21,790 | 490 |
| webserver #3 (subnet in us-west-2b): | 0 | 410 |
| webserver #4 (subnet in us-west-2b): | 0 | 428 |

Which recommendations can help ensure load-testing HTTP requests are evenly distributed across the four web servers? Choose 2 answers

- A. Launch and run the load-tester EC2 instance from us-east-1 instead.
- B. Re-configure the load-testing software to re-resolve DNS for each web request.
- C. Use a 3rd-party load-testing service which offers globally-distributed test clients.
- D. Configure ELB and Auto Scaling to distribute across us-west-2a and us-west-2c.
- E. Configure ELB session stickiness to use the app-specific session cookie

Answer: BE

NEW QUESTION 219

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 logged 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 state

Answer: D

NEW QUESTION 223

Which of the following are valid SNS delivery transports? Choose 2 answers

- A. HTTP
- B. UDP
- C. SNIS
- D. DynamoDB
- E. Named Pipes

Answer: AC

NEW QUESTION 228

Which DynamoDB limits can be raised by contacting AWS support? Choose 2 answers

- A. The number of hash keys per account
- B. The maximum storage used per account
- C. The number of tables per account
- D. The number of local secondary indexes per account
- E. The number of provisioned throughput units per account

Answer: CE

NEW QUESTION 233

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 235

Company C is currently hosting their corporate site in an Amazon S3 bucket with Static Website Hosting enabled. Currently, when visitors go to <http://www.companyc.com> the index.html page is returned. Company C now would like a new page welcome.html to be returned when a visitor enters <http://www.companyc.com> in the browser.

Which of the following steps will allow Company C to meet this requirement? Choose 2 answers

- A. Upload an html page named welcome.html to their S3 bucket
- B. Create a welcome subfolder in their S3 bucket
- C. Set the Index Document property to welcome.html
- D. Move the index.html page to a welcome subfolder
- E. Set the Error Document property to welcome.html

Answer: AC

NEW QUESTION 240

Which of the following are valid arguments for an SNS Publish request? Choose 3 answers

- A. TopicArn
- B. Subject
- C. Destination
- D. Format
- E. Message
- F. Language

Answer: ABE

NEW QUESTION 241

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 244

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 246

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 251

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 252

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 257

You are inserting 1000 new items every second in a DynamoDB table. Once an hour these items are analyzed and then are no longer needed. You need to minimize provisioned throughput, storage, and API calls.

Given these requirements, what is the most efficient way to manage these Items after the analysis?

- A. Retain the items in a single table
- B. Delete items indMdually over a 24 hour period
- C. Delete the table and create a new table per hour
- D. Create a new table per hour

Answer: C

NEW QUESTION 259

Which features can be used to restrict access to data in S3? Choose 2 answers

- A. Use S3 Virtual Hosting
- B. Set an S3 Bucket policy.
- C. Enable IAM Identity Federation.
- D. Set an S3 ACL on the bucket or the object.
- E. Create a C|oudFront distribution for the bucket

Answer: CD

NEW QUESTION 261

Company B provides an online image recognition service and utilizes SOS 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 264

What is the maximum number of S3 Buckets available per AWS account?

- A. 100 per region
- B. there is no limit
- C. 100 per account
- D. 500 per account
- E. 100 per IAM user

Answer: C

NEW QUESTION 268

What is the format of structured notification messages sent by Amazon SNS?

- A. An XML object containing MessageId, UnsubscribeURL, Subject, IVmessage and other values
- B. An JSON object containing MessageId, DuplicateFlag, IVmessage and other values
- C. An XML object containing MessageId, DuplicateFlag, IVmessage and other values
- D. An JSON object containing MessageId, unsubscribeURL, Subject, IVmessage and other values

Answer: D

NEW QUESTION 270

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 275

An Amazon S3 bucket, "myawsbucket" is configured with website hosting in Tokyo region, what is the region-specific website endpoint?

- A. www.myawsbucket.ap-northeast-1.amazonaws.com
- B. myawsbucket.s3-website-ap-northeast-1.amazonaws.com
- C. myawsbucket.amazonaws.com
- D. myawsbucket.tokyo.amazonaws.com

Answer: B

NEW QUESTION 277

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 282

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 286

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 291

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 294

Which of the following statements about SQS is true?

- A. Messages will be delivered exactly once and messages will be delivered in First in, First out order
- B. Messages will be delivered exactly once and message delivery order is indeterminate
- C. Messages will be delivered one or more times and messages will be delivered in First in, First out order
- D. Messages will be delivered one or more times and message delivery order is indeterminate

Answer: D

NEW QUESTION 299

.....

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