

## AWS-Solution-Architect-Associate Dumps

### Amazon AWS Certified Solutions Architect - Associate

<https://www.certleader.com/AWS-Solution-Architect-Associate-dumps.html>



#### NEW QUESTION 1

You are trying to launch an EC2 instance, however the instance seems to go into a terminated status immediately. What would probably not be a reason that this is happening?

- A. The AMI is missing a required part.
- B. The snapshot is corrupt.
- C. You need to create storage in EBS first.
- D. You've reached your volume limit

**Answer: C**

#### Explanation:

Amazon EC2 provides a virtual computing environments, known as an instance.

After you launch an instance, AWS recommends that you check its status to confirm that it goes from the pending status to the running status, the not terminated status.

The following are a few reasons why an Amazon EBS-backed instance might immediately terminate: You've reached your volume limit.

The AMI is missing a required part. The snapshot is corrupt. Reference:

[http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using\\_InstanceStraightToTerminated.html](http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_InstanceStraightToTerminated.html)

#### NEW QUESTION 2

In the context of AWS support, why must an EC2 instance be unreachable for 20 minutes rather than allowing customers to open tickets immediately?

- A. Because most reachability issues are resolved by automated processes in less than 20 minutes
- B. Because all EC2 instances are unreachable for 20 minutes every day when AWS does routine maintenance
- C. Because all EC2 instances are unreachable for 20 minutes when first launched
- D. Because of all the reasons listed here

**Answer: A**

#### Explanation:

An EC2 instance must be unreachable for 20 minutes before opening a ticket, because most reachability issues are resolved by automated processes in less than 20 minutes and will not require any action on the part of the customer. If the instance is still unreachable after this time frame has passed, then you should open a case with support.

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

#### NEW QUESTION 3

Does DynamoDB support in-place atomic updates?

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

**Answer: A**

#### Explanation:

DynamoDB supports in-place atomic updates.

Reference:

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

#### NEW QUESTION 4

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

**Answer: D**

#### 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 5

You need to migrate a large amount of data into the cloud that you have stored on a hard disk and you decide that the best way to accomplish this is with AWS Import/Export and you mail the hard disk to AWS. Which of the following statements is incorrect in regards to AWS Import/Export?

- A. It can export from Amazon S3
- B. It can Import to Amazon Glacier
- C. It can export from Amazon Glacier.
- D. It can Import to Amazon EBS

**Answer:**

C

**Explanation:**

AWS Import/Export supports: Import to Amazon S3

Export from Amazon S3 Import to Amazon EBS Import to Amazon Glacier

AWS Import/Export does not currently support export from Amazon EBS or Amazon Glacier. Reference:

<https://docs.aws.amazon.com/AWSImportExport/latest/DG/whatisdisk.html>

**NEW QUESTION 6**

You are checking the workload on some of your General Purpose (SSD) and Provisioned IOPS (SSD) volumes and it seems that the I/O latency is higher than you require. You should probably check the to make sure that your application is not trying to drive more IOPS than you have provisioned.

- A. Amount of IOPS that are available
- B. Acknowledgement from the storage subsystem
- C. Average queue length
- D. Time it takes for the I/O operation to complete

**Answer: C****Explanation:**

In EBS workload demand plays an important role in getting the most out of your General Purpose (SSD) and Provisioned IOPS (SSD) volumes. In order for your volumes to deliver the amount of IOPS that are available, they need to have enough I/O requests sent to them. There is a relationship between the demand on the volumes, the amount of IOPS that are available to them, and the latency of the request (the amount of time it takes for the I/O operation to complete).

Latency is the true end-to-end client time of an I/O operation; in other words, when the client sends a IO, how long does it take to get an acknowledgement from the storage subsystem that the IO read or write is complete.

If your I/O latency is higher than you require, check your average queue length to make sure that your application is not trying to drive more IOPS than you have provisioned. You can maintain high IOPS while keeping latency down by maintaining a low average queue length (which is achieved by provisioning more IOPS for your volume).

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

**NEW QUESTION 7**

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

**Answer: B****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 8**

You log in to IAM on your AWS console and notice the following message. "Delete your root access keys." Why do you think IAM is requesting this?

- A. Because the root access keys will expire as soon as you log out.
- B. Because the root access keys expire after 1 week.
- C. Because the root access keys are the same for all users.
- D. Because they provide unrestricted access to your AWS resource

**Answer: D****Explanation:**

In AWS an access key is required in order to sign requests that you make using the command-line interface (CLI), using the AWS SDKs, or using direct API calls. Anyone who has the access key for your root account has unrestricted access to all the resources in your account, including billing information. One of the best ways to protect your account is to not have an access key for your root account. We recommend that unless you must have a root access key (this is very rare), that you do not generate one. Instead, AWS best practice is to create one or more AWS Identity and Access Management (IAM) users, give them the necessary permissions, and use IAM users for everyday interaction with AWS.

Reference:

<http://docs.aws.amazon.com/general/latest/gr/aws-access-keys-best-practices.html#root-password>

**NEW QUESTION 9**

You are setting up a very complex financial services grid and so far it has 5 Elastic IP (EIP) addresses.

You go to assign another EIP address, but all accounts are limited to 5 Elastic IP addresses per region by default, so you aren't able to. What is the reason for this?

- A. For security reasons.
- B. Hardware restrictions.
- C. Public (IPv4) internet addresses are a scarce resource.
- D. There are only 5 network interfaces per instanc

**Answer: C****Explanation:**

Public (IPv4) internet addresses are a scarce resource. There is only a limited amount of public IP space available, and Amazon EC2 is committed to helping use that space efficiently.

By default, all accounts are limited to 5 Elastic IP addresses per region. If you need more than 5 Elastic IP addresses, AWS asks that you apply for your limit to be raised. They will ask you to think through your use case and help them understand your need for additional addresses.

Reference: [http://aws.amazon.com/ec2/faqs/#How\\_many\\_instances\\_can\\_I\\_run\\_in\\_Amazon\\_EC2](http://aws.amazon.com/ec2/faqs/#How_many_instances_can_I_run_in_Amazon_EC2)

**NEW QUESTION 10**

You need to measure the performance of your EBS volumes as they seem to be under performing. You have come up with a measurement of 1,024 KB I/O but your colleague tells you that EBS volume performance is measured in IOPS. How many IOPS is equal to 1,024 KB I/O?

- A. 16
- B. 256
- C. 8
- D. 4

**Answer:** D

**Explanation:**

Several factors can affect the performance of Amazon EBS volumes, such as instance configuration, I/O characteristics, workload demand, and storage configuration.

IOPS are input/output operations per second. Amazon EBS measures each I/O operation per second (that is 256 KB or smaller) as one IOPS. I/O operations that are larger than 256 KB are counted in 256 KB capacity units.

For example, a 1,024 KB I/O operation would count as 4 IOPS.

When you provision a 4,000 IOPS volume and attach it to an EBS-optimized instance that can provide the necessary bandwidth, you can transfer up to 4,000 chunks of data per second (provided that the I/O does not exceed the 128 MB/s per volume throughput limit of General Purpose (SSD) and Provisioned IOPS (SSD) volumes).

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

**NEW QUESTION 10**

Having set up a website to automatically be redirected to a backup website if it fails, you realize that there are different types of failovers that are possible. You need all your resources to be available the majority of the time. Using Amazon Route 53 which configuration would best suit this requirement?

- A. Active-active failover.
- B. Non
- C. Route 53 can't failover.
- D. Active-passive failover.
- E. Active-active-passive and other mixed configuration

**Answer:** A

**Explanation:**

You can set up a variety of failover configurations using Amazon Route 53 alias: weighted, latency, geolocation routing, and failover resource record sets.

Active-active failover: Use this failover configuration when you want all of your resources to be available the majority of the time. When a resource becomes unavailable, Amazon Route 53 can detect that it's unhealthy and stop including it when responding to queries.

Active-passive failover: Use this failover configuration when you want a primary group of resources to be available the majority of the time and you want a secondary group of resources to be on standby in case all of the primary resources become unavailable. When responding to queries, Amazon Route 53 includes only the healthy primary resources. If all of the primary resources are unhealthy, Amazon Route 53 begins to include only the healthy secondary resources in response to DNS queries.

Active-active-passive and other mixed configurations: You can combine alias and non-alias resource record sets to produce a variety of Amazon Route 53 behaviors.

Reference: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html>

**NEW QUESTION 11**

True or False: In Amazon Route 53, you can create a hosted zone for a top-level domain (TLD).

- A. FALSE
- B. False, Amazon Route 53 automatically creates it for you.
- C. True, only if you send an XML document with a CreateHostedZoneRequest element for TLD.
- D. TRUE

**Answer:** A

**Explanation:**

In Amazon Route 53, you cannot create a hosted zone for a top-level domain (TLD).

Reference: [http://docs.aws.amazon.com/Route53/latest/APIReference/API\\_CreateHostedZone.html](http://docs.aws.amazon.com/Route53/latest/APIReference/API_CreateHostedZone.html)

**NEW QUESTION 16**

Which of the following statements is true of creating a launch configuration using an EC2 instance?

- A. The launch configuration can be created only using the Query APIs.
- B. Auto Scaling automatically creates a launch configuration directly from an EC2 instance.
- C. A user should manually create a launch configuration before creating an Auto Scaling group.
- D. The launch configuration should be created manually from the AWS CL

**Answer:** B

**Explanation:**

You can create an Auto Scaling group directly from an EC2 instance. When you use this feature, Auto Scaling automatically creates a launch configuration for you as well.

Reference:

<http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/create-ic-with-instanceID.html>

#### NEW QUESTION 19

You have been using T2 instances as your CPU requirements have not been that intensive. However you now start to think about larger instance types and start looking at M and IV|3 instances. You are a little confused as to the differences between them as they both seem to have the same ratio of CPU and memory. Which statement below is incorrect as to why you would use one over the other?

- A. M3 instances are less expensive than M1 instances.
- B. IV|3 instances are configured with more swap memory than M instances.
- C. IV|3 instances provide better, more consistent performance than M instances for most use-cases.
- D. M3 instances also offer SSD-based instance storage that delivers higher I/O performance.

**Answer: B**

#### Explanation:

Amazon EC2 allows you to set up and configure everything about your instances from your operating system up to your applications. An Amazon Machine Image (AMI) is simply a packaged-up environment that includes all the necessary bits to set up and boot your instance.

M1 and M3 Standard instances have the same ratio of CPU and memory, some reasons below as to why you would use one over the other.

IV|3 instances provide better, more consistent performance than M instances for most use-cases. M3 instances also offer SSD-based instance storage that delivers higher I/O performance.

M3 instances are also less expensive than M1 instances. Due to these reasons, we recommend M3 for applications that require general purpose instances with a balance of compute, memory, and network resources.

However, if you need more disk storage than what is provided in M3 instances, you may still find M1 instances useful for running your applications.

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

#### NEW QUESTION 22

A user wants to achieve High Availability with PostgreSQL DB. Which of the below mentioned functionalities helps achieve HA?

- A. Multi AZ
- B. Read Replica
- C. Multi region
- D. PostgreSQL does not support HA

**Answer: A**

#### Explanation:

The Multi AZ feature allows the user to achieve High Availability. For Multi AZ, Amazon RDS automatically provisions and maintains a synchronous "standby" replica in a different Availability Zone. Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html>

#### NEW QUESTION 27

After setting up several database instances in Amazon Relational Database Service (Amazon RDS) you decide that you need to track the performance and health of your databases. How can you do this?

- A. Subscribe to Amazon RDS events to be notified when changes occur with a DB instance, DB snapshot, DB parameter group, or DB security group.
- B. Use the free Amazon CloudWatch service to monitor the performance and health of a DB instance.
- C. All of the items listed will track the performance and health of a database.
- D. View, download, or watch database log files using the Amazon RDS console or Amazon RDS API
- E. You can also query some database log files that are loaded into database tables.

**Answer: C**

#### Explanation:

Amazon Relational Database Service (Amazon RDS) is a web service that makes it easier to set up, operate, and scale a relational database in the cloud. It provides cost-efficient, resizable capacity for an industry-standard relational database and manages common database administration tasks.

There are several ways you can track the performance and health of a database or a DB instance. You can:

Use the free Amazon CloudWatch service to monitor the performance and health of a DB instance. Subscribe to Amazon RDS events to be notified when changes occur with a DB instance, DB snapshot, DB parameter group, or DB security group.

View, download, or watch database log files using the Amazon RDS console or Amazon RDS APIs. You can also query some database log files that are loaded into database tables.

Use the AWS CloudTrail service to record AWS calls made by your AWS account. The calls are recorded in log files and stored in an Amazon S3 bucket.

Reference: [http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Monitoring.html](http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Monitoring.html)

#### NEW QUESTION 31

You are building a system to distribute confidential documents to employees. Using CloudFront, what method could be used to serve content that is stored in S3, but not publically accessible from S3 directly?

- A. Add the CloudFront account security group "amazon-cf/amazon-cf-sg" to the appropriate S3 bucket policy.
- B. Create a S3 bucket policy that lists the CloudFront distribution ID as the Principal and the target bucket as the Amazon Resource Name (ARN).
- C. Create an Identity and Access Management (IAM) User for CloudFront and grant access to the objects in your S3 bucket to that IAM User.
- D. Create an Origin Access Identity (OAI) for CloudFront and grant access to the objects in your S3 bucket to that OAI.

**Answer: D**

#### Explanation:

You restrict access to Amazon S3 content by creating an origin access identity, which is a special CloudFront user. You change Amazon S3 permissions to give the origin access identity permission to access your objects, and to remove permissions from everyone else. When your users access your Amazon S3 objects using CloudFront URLs, the CloudFront origin access identity gets the objects on your users' behalf. If your users try to access objects using Amazon S3 URLs, they're denied access. The origin access identity has permission to access objects in your Amazon S3 bucket, but users don't. Reference:

<http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html>

**NEW QUESTION 36**

A user has attached 1 EBS volume to a VPC instance. The user wants to achieve the best fault tolerance of data possible. Which of the below mentioned options can help achieve fault tolerance?

- A. Attach one more volume with RAID 1 configuration.
- B. Attach one more volume with RAID 0 configuration.
- C. Connect multiple volumes and stripe them with RAID 6 configuration.
- D. Use the EBS volume as a root device

**Answer:** A

**Explanation:**

The user can join multiple provisioned IOPS volumes together in a RAID 1 configuration to achieve better fault tolerance. RAID 1 does not provide a write performance improvement; it requires more bandwidth than non-RAID configurations since the data is written simultaneously to multiple volumes.  
Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html>

**NEW QUESTION 39**

Select a true statement about Amazon EC2 Security Groups (EC2-Classic).

- A. After you launch an instance in EC2-Classic, you can't change its security groups.
- B. After you launch an instance in EC2-Classic, you can change its security groups only once.
- C. After you launch an instance in EC2-Classic, you can only add rules to a security group.
- D. After you launch an instance in EC2-Classic, you cannot add or remove rules from a security group

**Answer:** A

**Explanation:**

After you launch an instance in EC2-Classic, you can't change its security groups. However, you can add rules to or remove rules from a security group, and those changes are automatically applied to all instances that are associated with the security group.  
Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-network-security.html>

**NEW QUESTION 44**

Which one of the following answers is not a possible state of Amazon CloudWatch Alarm?

- A. INSUFFICIENT\_DATA
- B. ALARM
- C. OK
- D. STATUS\_CHECK\_FAILED

**Answer:** D

**Explanation:**

Amazon CloudWatch Alarms have three possible states: OK: The metric is within the defined threshold ALARM: The metric is outside of the defined threshold INSUFFICIENT\_DATA: The alarm has just started, the metric is not available, or not enough data is available for the metric to determine the alarm state  
Reference: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/AlarmThatSendsEmail.html>

**NEW QUESTION 48**

A user is planning to launch a scalable web application. Which of the below mentioned options will not affect the latency of the application?

- A. Region.
- B. Provisioned IOPS.
- C. Availability Zone.
- D. Instance size

**Answer:** C

**Explanation:**

In AWS, the instance size decides the I/O characteristics. The provisioned IOPS ensures higher throughput, and lower latency. The region does affect the latency; latency will always be less when the instance is near to the end user. Within a region the user uses any AZ and this does not affect the latency. The AZ is mainly for fault tolerance or HA.  
Reference: [http://media.amazonwebservices.com/AWS\\_Security\\_Best\\_Practices.pdf](http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf)

**NEW QUESTION 49**

An organization has a statutory requirement to protect the data at rest for data stored in EBS volumes. Which of the below mentioned options can the organization use to achieve data protection?

- A. Data replication.
- B. Data encryption.
- C. Data snapshot.
- D. All the options listed here

**Answer:** D

**Explanation:**

For protecting the Amazon EBS data at REST, the user can use options, such as Data Encryption (Windows / Linux / third party based), Data Replication (AWS

internally replicates data for redundancy),  
and Data Snapshot (for point in time backup).  
Reference: [http://media.amazonwebservices.com/AWS\\_Security\\_Best\\_Practices.pdf](http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf)

**NEW QUESTION 54**

A user is running a batch process which runs for 1 hour every day. Which of the below mentioned options is the right instance type and costing model in this case if the user performs the same task for the whole year?

- A. EBS backed instance with on-demand instance pricing.
- B. EBS backed instance with heavy utilized reserved instance pricing.
- C. EBS backed instance with low utilized reserved instance pricing.
- D. Instance store backed instance with spot instance pricing.

**Answer:** A

**Explanation:**

For Amazon Web Services, the reserved instance helps the user save money if the user is going to run the same instance for a longer period. Generally if the user uses the instances around 30-40% annually it is recommended to use RI. Here as the instance runs only for 1 hour daily it is not recommended to have RI as it will be costlier. The user should use on-demand with EBS in this case.

Reference: <http://aws.amazon.com/ec2/purchasing-options/reserved-instances/>

**NEW QUESTION 55**

You have just set up a large site for a client which involved a huge database which you set up with Amazon RDS to run as a Multi-AZ deployment. You now start to worry about what will happen if the database instance fails. Which statement best describes how this database will function if there is a database failure?

- A. Updates to your DB Instance are synchronously replicated across Availability Zones to the standby in order to keep both in sync and protect your latest database updates against DB Instance failure.
- B. Your database will not resume operation without manual administrative intervention.
- C. Updates to your DB Instance are asynchronously replicated across Availability Zones to the standby in order to keep both in sync and protect your latest database updates against DB Instance failure.
- D. Updates to your DB Instance are synchronously replicated across S3 to the standby in order to keep both in sync and protect your latest database updates against DB Instance failure.

**Answer:** A

**Explanation:**

Amazon Relational Database Service (Amazon RDS) is a managed service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity, while managing time-consuming database administration tasks, freeing you up to focus on your applications and business.

When you create or modify your DB Instance to run as a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous "standby" replica in a different Availability Zone. Updates to your DB Instance are synchronously replicated across Availability Zones to the standby in order to keep both in sync and protect your latest database updates against DB Instance failure.

During certain types of planned maintenance, or in the unlikely event of DB Instance failure or Availability Zone failure, Amazon RDS will automatically failover to the standby so that you can resume database writes and reads as soon as the standby is promoted. Since the name record for your DB Instance remains the same, your application can resume database operation without the need for manual administrative intervention. With Multi-AZ deployments, replication is transparent: you do not interact directly with the standby, and it cannot be used to serve read traffic. If you are using Amazon RDS for MySQL and are looking to scale read traffic beyond the capacity constraints of a single DB Instance, you can deploy one or more Read Replicas.

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

**NEW QUESTION 59**

Which IAM role do you use to grant AWS Lambda permission to access a DynamoDB Stream?

- A. Dynamic role
- B. Invocation role
- C. Execution role
- D. Event Source role

**Answer:** C

**Explanation:**

You grant AWS Lambda permission to access a DynamoDB Stream using an IAM role known as the "execution role".

Reference: <http://docs.aws.amazon.com/lambda/latest/dg/intro-permission-model.html>

**NEW QUESTION 64**

Name the disk storage supported by Amazon Elastic Compute Cloud (EC2).

- A. None of these
- B. Amazon AppStream store
- C. Amazon SNS store
- D. Amazon Instance Store

**Answer:** D

**Explanation:**

Amazon EC2 supports the following storage options: Amazon Elastic Block Store (Amazon EBS) Amazon EC2 Instance Store Amazon Simple Storage Service (Amazon S3)

Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Storage.html>

**NEW QUESTION 67**

A user is observing the EC2 CPU utilization metric on CloudWatch. The user has observed some interesting patterns while filtering over the 1 week period for a particular hour. The user wants to zoom that data point to a more granular period. How can the user do that easily with CloudWatch?

- A. The user can zoom a particular period by selecting that period with the mouse and then releasing the mouse
- B. The user can zoom a particular period by specifying the aggregation data for that period
- C. The user can zoom a particular period by double clicking on that period with the mouse
- D. The user can zoom a particular period by specifying the period in the Time Range

**Answer:** A

**Explanation:**

Amazon CloudWatch provides the functionality to graph the metric data generated either by the AWS services or the custom metric to make it easier for the user to analyse. The AWS CloudWatch console provides the option to change the granularity of a graph and zoom in to see data over a shorter time period. To zoom, the user has to click in the graph details pane, drag on the graph area for selection, and then release the mouse button.

Reference: [http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/zoom\\_in\\_on\\_graph.html](http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/zoom_in_on_graph.html)

**NEW QUESTION 71**

You need to create an Amazon Machine Image (AMI) for a customer for an application which does not appear to be part of the standard AWS AMI template that you can see in the AWS console. What are the alternative possibilities for creating an AMI on AWS?

- A. You can purchase an AMIs from a third party but cannot create your own AMI.
- B. You can purchase an AMIs from a third party or can create your own AMI.
- C. Only AWS can create AMIs and you need to wait till it becomes available.
- D. Only AWS can create AMIs and you need to request them to create one for you

**Answer:** B

**Explanation:**

You can purchase an AMIs from a third party, including AMIs that come with service contracts from organizations such as Red Hat. You can also create an AMI and sell it to other Amazon EC2 users. After you create an AMI, you can keep it private so that only you can use it, or you can share it with a specified list of AWS accounts. You can also make your custom AMI public so that the community can use it. Building a safe, secure, usable AMI for public consumption is a fairly straightforward process, if you follow a few simple guidelines.

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

**NEW QUESTION 76**

You are very concerned about security on your network because you have multiple programmers testing APIs and SDKs and you have no idea what is happening. You think CloudTrail may help but are not sure what it does. Which of the following statements best describes the AWS service CloudTrail?

- A. With AWS CloudTrail you can get a history of AWS API calls and related events for your account.
- B. With AWS CloudTrail you can get a history of IAM users for your account.
- C. With AWS CloudTrail you can get a history of S3 logfiles for your account.
- D. With AWS CloudTrail you can get a history of CloudFormation JSON scripts used for your account

**Answer:** A

**Explanation:**

With AWS CloudTrail, you can get a history of AWS API calls for your account, including API calls made via the AWS Management Console, the AWS SDKs, the command line tools, and higher-level AWS services. You can also identify which users and accounts called AWS APIs for services that support CloudTrail, the source IP address the calls were made from, and when the calls occurred.

You can identify which users and accounts called AWS for services that support CloudTrail, the source IP address the calls were made from, and when the calls occurred. You can integrate CloudTrail into applications using the API, automate trail creation for your organization, check the status of your trails, and control how administrators turn CloudTrail logging on and off.

Reference: [http://docs.aws.amazon.com/awsccloudtrail/latest/userguide/what\\_is\\_cloud\\_trail\\_top\\_level.html](http://docs.aws.amazon.com/awsccloudtrail/latest/userguide/what_is_cloud_trail_top_level.html)

**NEW QUESTION 78**

A user has deployed an application on his private cloud. The user is using his own monitoring tool. He wants to configure it so that whenever there is an error, the monitoring tool will notify him via SMS. Which of the below mentioned AWS services will help in this scenario?

- A. AWS SES
- B. AWS SNS
- C. None because the user infrastructure is in the private cloud.
- D. AWS SMS

**Answer:** B

**Explanation:**

Amazon Simple Notification Service (Amazon SNS) is a fast, flexible, and fully managed push messaging service. Amazon SNS can be used to make push notifications to mobile devices. Amazon SNS can

deliver notifications by SMS text message or email to the Amazon Simple Queue Service (SQS) queues or to any HTTP endpoint. In this case user can use the SNS APIs to send SMS.

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

**NEW QUESTION 83**

A user is currently building a website which will require a large number of instances in six months, when a demonstration of the new site will be given upon launch. Which of the below mentioned options allows the user to procure the resources beforehand so that they need not worry about infrastructure availability during the demonstration?

- A. Procure all the instances as reserved instances beforehand.

- B. Launch all the instances as part of the cluster group to ensure resource availability.
- C. Pre-warm all the instances one month prior to ensure resource availability.
- D. Ask AWS now to procure the dedicated instances in 6 month

**Answer:** A

**Explanation:**

Amazon Web Services has massive hardware resources at its data centers, but they are finite. The best way for users to maximize their access to these resources is by reserving a portion of the computing capacity that they require. This can be done through reserved instances. With reserved instances, the user literally reserves the computing capacity in the Amazon Web Services cloud.

Reference: [http://media.amazonwebservices.com/AWS\\_Building\\_Fault\\_Tolerant\\_Applications.pdf](http://media.amazonwebservices.com/AWS_Building_Fault_Tolerant_Applications.pdf)

**NEW QUESTION 87**

You need to create a JSON-formatted text file for AWS CloudFormation. This is your first template and the only thing you know is that the templates include several major sections but there is only one that is required for it to work. What is the only section required?

- A. Mappings
- B. Outputs
- C. Resources
- D. Conditions

**Answer:** C

**Explanation:**

AWS CloudFormation is a service that helps you model and set up your Amazon Web Services resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS. You create a template that describes all the AWS resources that you want (like Amazon EC2 instances or Amazon RDS DB instances), and AWS CloudFormation takes care of provisioning and configuring those resources for you.

A template is a JSON-formatted text file that describes your AWS infrastructure. Templates include several major sections.

The Resources section is the only section that is required.

The first character in the template must be an open brace ({}), and the last character must be a closed brace (}). The following template fragment shows the template structure and sections.

Reference: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-anatomy.html>

**NEW QUESTION 92**

You are planning and configuring some EBS volumes for an application. In order to get the most performance out of your EBS volumes, you should attach them to an instance with enough to support your volumes.

- A. Redundancy
- B. Storage
- C. Bandwidth
- D. Memory

**Answer:** C

**Explanation:**

When you plan and configure EBS volumes for your application, it is important to consider the configuration of the instances that you will attach the volumes to. In order to get the most performance out of your EBS volumes, you should attach them to an instance with enough bandwidth to support your volumes, such as an EBS-optimized instance or an instance with 10 Gigabit network connectivity. This is especially important when you use General Purpose (SSD) or Provisioned IOPS (SSD) volumes, or when you stripe multiple volumes together in a RAID configuration.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-ec2-config.html>

**NEW QUESTION 95**

Can a single EBS volume be attached to multiple EC2 instances at the same time?

- A. Yes
- B. No
- C. Only for high-performance EBS volumes.
- D. Only when the instances are located in the US region

**Answer:** B

**Explanation:**

You can't attach an EBS volume to multiple EC2 instances. This is because it is equivalent to using a single hard drive with many computers at the same time.

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

**NEW QUESTION 98**

Your organization is in the business of architecting complex transactional databases. For a variety of reasons, this has been done on EBS. What is AWS's recommendation for customers who have architected databases using EBS for backups?

- A. Backups to Amazon S3 be performed through the database management system.
- B. Backups to AWS Storage Gateway be performed through the database management system.
- C. If you take regular snapshots no further backups are required.
- D. Backups to Amazon Glacier be performed through the database management system.

**Answer:** A

**Explanation:**

Data stored in Amazon EBS volumes is redundantly stored in multiple physical locations as part of normal operation of those services and at no additional charge.

However, Amazon EBS replication is stored within the same availability zone, not across multiple zones; therefore, it is highly recommended that you conduct regular snapshots to Amazon S3 for long-term data durability.

For customers who have architected complex transactional databases using EBS, it is recommended that backups to Amazon S3 be performed through the database management system so that distributed transactions and logs can be checkpointed.

AWS does not perform backups of data that are maintained on virtual disks attached to running instances on Amazon EC2.

Reference: <http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf>

**NEW QUESTION 102**

You have three Amazon EC2 instances with Elastic IP addresses in the US East (Virginia) region, and you want to distribute requests across all three IPs evenly for users for whom US East (Virginia) is the appropriate region.

How many EC2 instances would be sufficient to distribute requests in other regions?

- A. 3
- B. 9
- C. 2
- D. 1

**Answer:** D

**Explanation:**

If your application is running on Amazon EC2 instances in two or more Amazon EC2 regions, and if you have more than one Amazon EC2 instance in one or more regions, you can use latency-based routing to route traffic to the correct region and then use weighted resource record sets to route traffic to instances within the region based on weights that you specify.

For example, suppose you have three Amazon EC2 instances with Elastic IP addresses in the US East (Virginia) region and you want to distribute requests across all three IPs evenly for users for whom US East (Virginia) is the appropriate region. Just one Amazon EC2 instance is sufficient in the other regions, although you can apply the same technique to many regions at once.

Reference: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/Tutorials.html>

**NEW QUESTION 106**

You have just finished setting up an advertisement server in which one of the obvious choices for a service was Amazon Elastic Map Reduce( EMR) and are now troubleshooting some weird cluster states that you are seeing. Which of the below is not an Amazon EMR cluster state?

- A. STARTING
- B. STOPPED
- C. RUNNING
- D. WAITING

**Answer:** B

**Explanation:**

Amazon Elastic Map Reduce (EMR) is a web service that enables businesses, researchers, data analysts, and developers to easily and cost-effectively process vast amounts of data.

Amazon EMR historically referred to an Amazon EMR cluster (and all processing steps assigned to it) as a "cluster". Every cluster has a unique identifier that starts with "j-".

The different cluster states of an Amazon EMR cluster are listed below. STARTING — The cluster provisions, starts, and configures EC2 instances.

BOOTSTRAPPING — Bootstrap actions are being executed on the cluster. RUNNING — A step for the cluster is currently being run.

WAITING — The cluster is currently active, but has no steps to run. TERMINATING - The cluster is in the process of shutting down. TERMINATED - The cluster was shut down without error. TERMINATED\_WITH\_ERRORS - The cluster was shut down with errors.

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

**NEW QUESTION 107**

Which of the following would you use to list your AWS Import/Export jobs?

- A. Amazon RDS
- B. AWS Import/Export Web Service Tool
- C. Amazon S3 REST API
- D. AWS Elastic Beanstalk

**Answer:** C

**Explanation:**

You can list AWS Import/Export jobs with the ListJobs command using the command line client or REST API.

Reference: <http://docs.aws.amazon.com/AWSImportExport/latest/DG/ListingYourJobs.html>

**NEW QUESTION 109**

How many types of block devices does Amazon EC2 support?

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

**Answer:** C

**Explanation:**

Amazon EC2 supports 2 types of block devices. Reference:

<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/block-device-mapping-concepts.html>

**NEW QUESTION 110**

You need to set up security for your VPC and you know that Amazon VPC provides two features that you can use to increase security for your VPC: Security groups and network access control lists (ACLs). You start to look into security groups first. Which statement below is incorrect in relation to security groups?

- A. Are stateful: Return traffic is automatically allowed, regardless of any rules.
- B. Evaluate all rules before deciding whether to allow traffic.
- C. Support allow rules and deny rules.
- D. Operate at the instance level (first layer of defense).

**Answer: C**

**Explanation:**

Amazon VPC provides two features that you can use to increase security for your VPC:

Security groups—Act as a firewall for associated Amazon EC2 instances, controlling both inbound and outbound traffic at the instance level and supports allow rules only.

Network access control lists (ACLs)—Act as a firewall for associated subnets, controlling both inbound and outbound traffic at the subnet level and supports allow rules and deny rules.

Reference: [http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC\\_Security.html](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Security.html)

**NEW QUESTION 112**

You are setting up some IAM user policies and have also become aware that some services support resource-based permissions, which let you attach policies to the service's resources instead of to IAM users or groups. Which of the below statements is true in regards to resource-level permissions?

- A. All services support resource-level permissions for all actions.
- B. Resource-level permissions are supported by Amazon CloudFront
- C. All services support resource-level permissions only for some actions.
- D. Some services support resource-level permissions only for some action

**Answer: D**

**Explanation:**

AWS Identity and Access Management is a web service that enables Amazon Web Services (AWS) customers to manage users and user permissions in AWS. The service is targeted at organizations with multiple users or systems that use AWS products such as Amazon EC2, Amazon RDS, and the AWS Management Console. With IAM, you can centrally manage users, security credentials such as access keys, and permissions that control which AWS resources users can access.

In addition to supporting IAM user policies, some services support resource-based permissions, which let you attach policies to the service's resources instead of to IAM users or groups. Resource-based permissions are supported by Amazon S3, Amazon SNS, and Amazon SQS.

The resource-level permissions service supports IAM policies in which you can specify individual resources using Amazon Resource Names (ARNs) in the policy's Resource element.

Some services support resource-level permissions only for some actions.

Reference: [http://docs.aws.amazon.com/IAM/latest/UserGuide/Using\\_SpecificProducts.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_SpecificProducts.html)

**NEW QUESTION 117**

Any person or application that interacts with AWS requires security credentials. AWS uses these credentials to identify who is making the call and whether to allow the requested access. You have just set up a VPC network for a client and you are now thinking about the best way to secure this network. You set up a security group called vpcsecuritygroup. Which following statement is true in respect to the initial settings that will be applied to this security group if you choose to use the default settings for this group?

- A. Allow all inbound traffic and allow no outbound traffic.
- B. Allow no inbound traffic and allow all outbound traffic.
- C. Allow inbound traffic on port 80 only and allow all outbound traffic.
- D. Allow all inbound traffic and allow all outbound traffi

**Answer: B**

**Explanation:**

Amazon VPC provides advanced security features such as security groups and network access control lists to enable inbound and outbound filtering at the instance level and subnet level.

AWS assigns each security group a unique ID in the form sg-xxxxxxx. The following are the initial settings for a security group that you create:

Allow no inbound traffic Allow all outbound traffic

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

**NEW QUESTION 120**

A for a VPC is a collection of subnets (typically private) that you may want to designate for your backend RDS DB Instances.

- A. DB Subnet Set
- B. RDS Subnet Group
- C. DB Subnet Group
- D. DB Subnet Collection

**Answer: C**

**Explanation:**

DB Subnet Groups are a set of subnets (one per Availability Zone of a particular region) designed for your DB instances that reside in a VPC. They make easy to manage Multi-AZ deployments as well as the conversion from a Single-AZ to a Mutli-AZ one.

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

**NEW QUESTION 123**

Amazon Elastic Load Balancing is used to manage traffic on a fleet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a

region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits. Which of the following is not an advantage of ELB over an on-premise load balancer?

- A. ELB uses a four-tier, key-based architecture for encryption.
- B. ELB offers clients a single point of contact, and can also serve as the first line of defense against attacks on your network.
- C. ELB takes over the encryption and decryption work from the Amazon EC2 instances and manages it centrally on the load balancer.
- D. ELB supports end-to-end traffic encryption using TLS (previously SSL) on those networks that use secure HTTP (HTTPS) connections.

**Answer:** A

**Explanation:**

Amazon Elastic Load Balancing is used to manage traffic on a fleet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits:  
Takes over the encryption and decryption work from the Amazon EC2 instances and manages it centrally on the load balancer  
Offers clients a single point of contact, and can also serve as the first line of defense against attacks on your network  
When used in an Amazon VPC, supports creation and management of security groups associated with your Elastic Load Balancing to provide additional networking and security options  
Supports end-to-end traffic encryption using TLS (previously SSL) on those networks that use secure HTTP (HTTPS) connections. When TLS is used, the TLS server certificate used to terminate client connections can be managed centrally on the load balancer, rather than on every individual instance. Reference: <http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf>

**NEW QUESTION 127**

An organization has a statutory requirement to protect the data at rest for the S3 objects. Which of the below mentioned options need not be enabled by the organization to achieve data security?

- A. MFA delete for S3 objects
- B. Client side encryption
- C. Bucket versioning
- D. Data replication

**Answer:** D

**Explanation:**

AWS S3 provides multiple options to achieve the protection of data at REST. The options include Permission (Policy), Encryption (Client and Server Side), Bucket Versioning and MFA based delete. The user can enable any of these options to achieve data protection. Data replication is an internal facility by AWS where S3 replicates each object across all the Availability Zones and the organization need not enable it in this case.  
Reference: [http://media.amazonwebservices.com/AWS\\_Security\\_Best\\_Practices.pdf](http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf)

**NEW QUESTION 131**

Content and Media Server is the latest requirement that you need to meet for a client.

The client has been very specific about his requirements such as low latency, high availability, durability, and access control. Potentially there will be millions of views on this server and because of "spiky" usage patterns, operations teams will need to provision static hardware, network, and management resources to support the maximum expected need. The Customer base will be initially low but is expected to grow and become more geographically distributed. Which of the following would be a good solution for content distribution?

- A. Amazon S3 as both the origin server and for caching
- B. AWS Storage Gateway as the origin server and Amazon EC2 for caching
- C. AWS CloudFront as both the origin server and for caching
- D. Amazon S3 as the origin server and Amazon CloudFront for caching

**Answer:** D

**Explanation:**

As your customer base grows and becomes more geographically distributed, using a high-performance edge cache like Amazon CloudFront can provide substantial improvements in latency, fault tolerance, and cost.  
By using Amazon S3 as the origin server for the Amazon CloudFront distribution, you gain the advantages of fast in-network data transfer rates, simple publishing/caching workflow, and a unified security framework.  
Amazon S3 and Amazon CloudFront can be configured by a web service, the AWS Management Console, or a host of third-party management tools.  
Reference: [http://media.amazonwebservices.com/architecturecenter/AWS\\_ac\\_ra\\_media\\_02.pdf](http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_media_02.pdf)

**NEW QUESTION 133**

You are setting up your first Amazon Virtual Private Cloud (Amazon VPC) network so you decide you should probably use the AWS Management Console and the VPC Wizard. Which of the following is not an option for network architectures after launching the "Start VPC Wizard" in Amazon VPC page on the AWS Management Console?

- A. VPC with a Single Public Subnet Only
- B. VPC with a Public Subnet Only and Hardware VPN Access
- C. VPC with Public and Private Subnets and Hardware VPN Access
- D. VPC with a Private Subnet Only and Hardware VPN Access

**Answer:** B

**Explanation:**

Amazon VPC enables you to build a virtual network in the AWS cloud - no VPNs, hardware, or physical datacenters required.  
Your AWS resources are automatically provisioned in a ready-to-use default VPC. You can choose to create additional VPCs by going to Amazon VPC page on the AWS Management Console and click on the "Start VPC Wizard" button.  
You'll be presented with four basic options for network architectures. After selecting an option, you can modify the size and IP address range of the VPC and its subnets. If you select an option with Hardware VPN Access, you will need to specify the IP address of the VPN hardware on your network. You can modify the VPC to add more subnets or add or remove gateways at any time after the VPC has been created.

The four options are:

VPC with a Single Public Subnet Only VPC with Public and Private Subnets

VPC with Public and Private Subnets and Hardware VPN Access VPC with a Private Subnet Only and Hardware VPN Access Reference:

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

**NEW QUESTION 137**

Which one of the below doesn't affect Amazon CloudFront billing?

- A. Distribution Type
- B. Data Transfer Out
- C. Dedicated IP SSL Certificates
- D. Requests

**Answer:** A

**Explanation:**

Amazon CloudFront is a web service for content delivery. CloudFront delivers your content using a global network of edge locations and works seamlessly with Amazon S3 which durably stores the original and definitive versions of your files.

Amazon CloudFront billing is mainly affected by Data Transfer Out

Edge Location Traffic Distribution Requests

Dedicated IP SSL Certificates

Reference: <http://calculator.s3.amazonaws.com/index.html>

**NEW QUESTION 140**

Your company has multiple IT departments, each with their own VPC. Some VPCs are located within the same AWS account, and others in a different AWS account. You want to peer together all VPCs to enable the IT departments to have full access to each others' resources. There are certain limitations placed on VPC peering. Which of the following statements is incorrect in relation to VPC peering?

- A. Private DNS values cannot be resolved between instances in peered VPCs.
- B. You can have up to 3 VPC peering connections between the same two VPCs at the same time.
- C. You cannot create a VPC peering connection between VPCs in different regions.
- D. You have a limit on the number active and pending VPC peering connections that you can have per VPC.

**Answer:** B

**Explanation:**

To create a VPC peering connection with another VPC, you need to be aware of the following limitations and rules:

You cannot create a VPC peering connection between VPCs that have matching or overlapping CIDR blocks.

You cannot create a VPC peering connection between VPCs in different regions.

You have a limit on the number active and pending VPC peering connections that you can have per VPC. VPC peering does not support transitive peering relationships; in a VPC peering connection, your VPC will not have access to any other VPCs that the peer VPC may be peered with. This includes VPC peering connections that are established entirely within your own AWS account.

You cannot have more than one VPC peering connection between the same two VPCs at the same time. The Maximum Transmission Unit (MTU) across a VPC peering connection is 1500 bytes.

A placement group can span peered VPCs; however, you will not get full-bisection bandwidth between instances in peered VPCs.

Unicast reverse path forwarding in VPC peering connections is not supported.

You cannot reference a security group from the peer VPC as a source or destination for ingress or egress rules in your security group. Instead, reference CIDR blocks of the peer VPC as the source or destination of your security group's ingress or egress rules.

Private DNS values cannot be resolved between instances in peered VPCs. Reference:

<http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/vpc-peering-overview.html#vpc-peering-limitations>

**NEW QUESTION 141**

In the most recent company meeting, your CEO focused on the fact that everyone in the organization needs to make sure that all of the infrastructure that is built is truly scalable. Which of the following statements is incorrect in reference to scalable architecture?

- A. A scalable service is capable of handling heterogeneity.
- B. A scalable service is resilient.
- C. A scalable architecture won't be cost effective as it grows.
- D. Increasing resources results in a proportional increase in performance.

**Answer:** C

**Explanation:**

In AWS it is critical to build a scalable architecture in order to take advantage of a scalable infrastructure. The cloud is designed to provide conceptually infinite scalability. However, you cannot leverage all that scalability in infrastructure if your architecture is not scalable. Both have to work together. You will have to identify the monolithic components and bottlenecks in your architecture, identify the areas where you cannot leverage the on-demand provisioning capabilities in your architecture, and work to refactor your application, in order to leverage the scalable infrastructure and take advantage of the cloud.

Characteristics of a truly scalable application:

Increasing resources results in a proportional increase in performance A scalable service is capable of handling heterogeneity

A scalable service is operationally efficient A scalable service is resilient

A scalable service should become more cost effective when it grows (Cost per unit reduces as the number of units increases)

Reference: [http://media.amazonwebservices.com/AWS\\_Cloud\\_Best\\_Practices.pdf](http://media.amazonwebservices.com/AWS_Cloud_Best_Practices.pdf)

**NEW QUESTION 143**

You are playing around with setting up stacks using JSON templates in CloudFormation to try and understand them a little better. You have set up about 5 or 6 but now start to wonder if you are being charged for these stacks. What is AWS's billing policy regarding stack resources?

- A. You are not charged for the stack resources if they are not taking any traffic.
- B. You are charged for the stack resources for the time they were operating (even if you deleted the stack right away)

- C. You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 60 minutes)
- D. You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 30 minutes)

**Answer:** B

**Explanation:**

A stack is a collection of AWS resources that you can manage as a single unit. In other words, you can create, update, or delete a collection of resources by creating, updating, or deleting stacks. All the resources in a stack are defined by the stack's AWS CloudFormation template. A stack, for instance, can include all the resources required to run a web application, such as a web server, a database, and networking rules. If you no longer require that web application, you can simply delete the stack, and all of its related resources are deleted.

You are charged for the stack resources for the time they were operating (even if you deleted the stack right away).

Reference: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacks.html>

**NEW QUESTION 145**

All Amazon EC2 instances are assigned two IP addresses at launch. Which are those?

- A. 2 Elastic IP addresses
- B. A private IP address and an Elastic IP address
- C. A public IP address and an Elastic IP address
- D. A private IP address and a public IP address

**Answer:** D

**Explanation:**

In Amazon EC2-Classical every instance is given two IP Addresses: a private IP address and a public IP address

Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-instance-addressing.html#differences>

**NEW QUESTION 150**

While controlling access to Amazon EC2 resources, which of the following acts as a firewall that controls the traffic allowed to reach one or more instances?

- A. A security group
- B. An instance type
- C. A storage cluster
- D. An object

**Answer:** A

**Explanation:**

A security group acts as a firewall that controls the traffic allowed to reach one or more instances. When you launch an instance, you assign it one or more security groups.

Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/UsingIAM.html>

**NEW QUESTION 155**

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 Email Service.
- C. AWS NMobile Communication Service.
- D. AWS Simple Queue 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 157**

You have been asked to set up a database in AWS that will require frequent and granular updates. You know that you will require a reasonable amount of storage space but are not sure of the best option. What is the recommended storage option when you run a database on an instance with the above criteria?

- A. Amazon S3
- B. Amazon EBS
- C. AWS Storage Gateway
- D. Amazon Glacier

**Answer:** B

**Explanation:**

Amazon EBS provides durable, block-level storage volumes that you can attach to a running Amazon EC2 instance. You can use Amazon EBS as a primary storage device for data that requires frequent and granular updates. For example, Amazon EBS is the recommended storage option when you run a database on an instance.

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

**NEW QUESTION 159**

A user has hosted an application on EC2 instances. The EC2 instances are configured with ELB and Auto Scaling. The application server session time out is 2 hours. The user wants to configure connection draining to ensure that all in-flight requests are supported by ELB even though the instance is being deregistered. What time out period should the user specify for connection draining?

- A. 1 hour
- B. 30 minutes
- C. 5 minutes
- D. 2 hours

**Answer:** A

**Explanation:**

The Elastic Load Balancer connection draining feature causes the load balancer to stop sending new requests to the back-end instances when the instances are deregistering or become unhealthy, while ensuring that in-flight requests continue to be served. The user can specify a maximum time of 3600 seconds (1 hour) for the load balancer to keep the connections alive before reporting the instance as deregistered. If the user does not specify the maximum timeout period, by default, the load balancer will close the connections to the deregistering instance after 300 seconds.

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/config-conn-drain.html>

**NEW QUESTION 161**

Identify a true statement about the On-Demand instances purchasing option provided by Amazon EC2.

- A. Pay for the instances that you use by the hour, with no long-term commitments or up-front payments.
- B. Make a low, one-time, up-front payment for an instance, reserve it for a one- or three-year term, and pay a significantly lower hourly rate for these instances.
- C. Pay for the instances that you use by the hour, with long-term commitments or up-front payments.
- D. Make a high, one-time, all-front payment for an instance, reserve it for a one- or three-year term, and pay a significantly higher hourly rate for these instance

**Answer:** A

**Explanation:**

On-Demand instances allow you to pay for the instances that you use by the hour, with no long-term commitments or up-front payments.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/reserved-instances-offerings.html>

**NEW QUESTION 166**

You have a Business support plan with AWS. One of your EC2 instances is running Microsoft Windows Server 2008 R2 and you are having problems with the software. Can you receive support from AWS for this software?

- A. Yes
- B. No, AWS does not support any third-party software.
- C. No, Microsoft Windows Server 2008 R2 is not supported.
- D. No, you need to be on the enterprise support plan

**Answer:** A

**Explanation:**

Third-party software support is available only to AWS Support customers enrolled for Business or Enterprise Support. Third-party support applies only to software running on Amazon EC2 and does not extend to assisting with on-premises software. An exception to this is a VPN tunnel configuration running supported devices for Amazon VPC.

Reference: <https://aws.amazon.com/premiumsupport/features/>

**NEW QUESTION 171**

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

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

**Answer:** D

**Explanation:**

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

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

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

**NEW QUESTION 172**

You need to set up a security certificate for a client's e-commerce website as it will use the HTTPS protocol. Which of the below AWS services do you need to access to manage your SSL server certificate?

- A. AWS Directory Service
- B. AWS Identity & Access Management
- C. AWS CloudFormation

D. Amazon Route 53

**Answer:** B

**Explanation:**

AWS Identity and Access Management (IAM) is a web service that enables Amazon Web Services (AWS) customers to manage users and user permissions in AWS.

All your SSL server certificates are managed by AWS Identity and Access management (IAM). Reference:  
<http://docs.aws.amazon.com/IAM/latest/UserGuide/ManagingServerCerts.htm>

**NEW QUESTION 174**

A customer has a 10 GB AWS Direct Connect connection to an AWS region where they have a web application hosted on Amazon Elastic Computer Cloud (EC2). The application has dependencies on an on-premises mainframe database that uses a BASE (Basic Available. Sort stale Eventual consistency) rather than an ACID (Atomicity. Consistency isolation. Durability) consistency model.

The application is exhibiting undesirable behavior because the database is not able to handle the volume of writes. How can you reduce the load on your on-premises database resources in the most cost-effective way?

- A. Use an Amazon Elastic Map Reduce (EMR) S3DistCp as a synchronization mechanism between the on-premises database and a Hadoop cluster on AWS.
- B. Modify the application to write to an Amazon SQS queue and develop a worker process to flush the queue to the on-premises database.
- C. Modify the application to use DynamoDB to feed an EMR cluster which uses a map function to write to the on-premises database.
- D. Provision an RDS read-replica database on AWS to handle the writes and synchronize the two databases using Data Pipeline.

**Answer:** A

**Explanation:**

Reference: <https://aws.amazon.com/blogs/aws/category/amazon-elastic-map-reduce/>

**NEW QUESTION 179**

You have launched an EC2 instance with four (4) 500GB EBS Provisioned IOPS volumes attached. The EC2 Instance is EBS-Optimized and supports 500 Mbps throughput between EC2 and EBS. The two EBS volumes are configured as a single RAID 0 device, and each Provisioned IOPS volume is provisioned with 4,000 IOPS (4,000 16KB reads or writes) for a total of 16,000 random IOPS on the instance. The EC2 Instance initially delivers the expected 16,000 IOPS random read and write performance. Sometime later in order to increase the total random I/O performance of the instance, you add an additional two 500 GB EBS Provisioned IOPS volumes to the RAID. Each volume is provisioned to 4,000 IOPS like the original four for a total of 24,000 IOPS on the EC2 instance. Monitoring shows that the EC2 instance CPU utilization increased from 50% to 70%, but the total random IOPS measured at the instance level does not increase at all. What is the problem and a valid solution?

- A. Larger storage volumes support higher Provisioned IOPS rates: increase the provisioned volumestorage of each of the 6 EBS volumes to 1TB
- B. The EBS-Optimized throughput limits the total IOPS that can be utilized. Use an EBS-Optimized instance that provides larger throughput.
- C. Small block sizes cause performance degradation, limiting the I/O throughput, configure the instance device driver and file system to use 64KB blocks to increase throughput.
- D. RAID 0 only scales linearly to about 4 devices, use RAID 0 with 4 EBS Provisioned IOPS volumes but increase each Provisioned IOPS EBS volume to 6,000 IOPS.
- E. The standard EBS instance root volume limits the total IOPS rate, change the instance root volume to also be a 500GB 4,000 Provisioned IOPS volume.

**Answer:** E

**NEW QUESTION 180**

You need a persistent and durable storage to trace call activity of an IVR (Interactive Voice Response) system. Call duration is mostly in the 2-3 minutes timeframe. Each traced call can be either active or terminated. An external application needs to know each minute the list of currently active calls, which are usually a few calls/second. Put once per month there is a periodic peak up to 1000 calls/second for a few hours. The system is open 24/7 and any downtime should be avoided.

Historical data is periodically archived to files. Cost saving is a priority for this project.

What database implementation would better fit this scenario, keeping costs as low as possible?

- A. Use RDS Multi-AZ with two tables, one for "Active calls" and one for "Terminated calls". In this way the "Active calls" table is always small and effective to access.
- B. Use DynamoDB with a "Calls" table and a Global Secondary Index on a "IsActive" attribute that is present for active calls only. In this way the Global Secondary index is sparse and more effective.
- C. Use DynamoDB with a "Calls" table and a Global secondary index on a "State" attribute that can equal to "active" or "terminated" in this way the Global Secondary index can be used for all items in the table.
- D. Use RDS Multi-AZ with a "CALLS" table and an Indexed "STATE" field that can be equal to "ACTIVE" or "TERMINATED". In this way the SQL query is optimized by the use of the Index.

**Answer:** A

**NEW QUESTION 182**

A web design company currently runs several FTP servers that their 250 customers use to upload and download large graphic files. They wish to move this system to AWS to make it more scalable, but they wish to maintain customer privacy and keep costs to a minimum.

What AWS architecture would you recommend?

- A. Ask their customers to use an S3 client instead of an FTP client.
- B. Create a single S3 bucket. Create an IAM user for each customer. Put the IAM Users in a Group that has an IAM policy that permits access to sub-directories within the bucket via use of the 'username' Policy variable.
- C. Create a single S3 bucket with Reduced Redundancy Storage turned on and ask their customers to use an S3 client instead of an FTP client. Create a bucket for each customer with a Bucket Policy that permits access only to that one customer.
- D. Create an auto-scaling group of FTP servers with a scaling policy to automatically scale-in when minimum network traffic on the auto-scaling group is below a given threshold.
- E. Load a central list of ftp users from S3 as part of the user Data startup script on each Instance.

F. Create a single S3 bucket with Requester Pays turned on and ask their customers to use an S3 client instead of an FTP client. Create a bucket for each customer with a Bucket Policy that permits access only to that one customer.

**Answer:** A

**NEW QUESTION 187**

You are responsible for a legacy web application whose server environment is approaching end of life. You would like to migrate this application to AWS as quickly as possible, since the application environment currently has the following limitations:

The VM's single 10GB VMDK is almost full. The virtual network interface still uses the 10Mbps driver, which leaves your 100Mbps WAN connection completely underutilized.

It is currently running on a highly customized Windows VM within a VMware environment. You do not have the installation media.

This is a mission-critical application with an RTO (Recovery Time Objective) of 8 hours. RPO (Recovery Point Objective) of 1 hour. How could you best migrate this application to AWS while meeting your business continuity requirements?

- A. Use the EC2 VM Import Connector for vCenter to import the VM into EC2.
- B. Use Import/Export to import the VM as an ESS snapshot and attach to EC2.
- C. Use S3 to create a backup of the VM and restore the data into EC2.
- D. Use the ec2-bundle-instance API to import an image of the VM into EC2.

**Answer:** A

**NEW QUESTION 191**

You have deployed a web application targeting a global audience across multiple AWS Regions under the domain name.example.com. You decide to use Route53 Latency-Based Routing to serve web requests to users from the region closest to the user. To provide business continuity in the event of server downtime, you configure weighted record sets associated with two web servers in separate Availability Zones per region. During a DR test, you notice that when you disable all web servers in one of the regions, Route53 does not automatically direct all users to the other region. What could be happening? (Choose 2 answers)

- A. Latency resource record sets cannot be used in combination with weighted resource record sets.
- B. You did not setup an HTTP health check for one or more of the weighted resource record sets associated with the disabled web servers.
- C. The value of the weight associated with the latency alias resource record set in the region with the disabled servers is higher than the weight for the other region.
- D. One of the two working web servers in the other region did not pass its HTTP health check.
- E. You did not set "Evaluate Target Health" to "Yes" on the latency alias resource record set associated with example.com in the region where you disabled the servers.

**Answer:** BE

**Explanation:**

How Health Checks Work in Complex Amazon Route 53 Configurations

Checking the health of resources in complex configurations works much the same way as in simple configurations. However, in complex configurations, you use a combination of alias resource record sets (including weighted alias, latency alias, and failover alias) and nonalias resource record sets to build a decision tree that gives you greater control over how Amazon Route 53 responds to requests.

For more information, see How Health Checks Work in Simple Amazon Route 53 Configurations.

For example, you might use latency alias resource record sets to select a region close to a user and use weighted resource record sets for two or more resources within each region to protect against the failure of a single endpoint or an Availability Zone. The following diagram shows this configuration.

Here's how Amazon EC2 and Amazon Route 53 are configured:

You have Amazon EC2 instances in two regions, us-east-1 and ap-southeast-2. You want Amazon Route 53 to respond to queries by using the resource record sets in the region that provides the lowest latency for your customers, so you create a latency alias resource record set for each region.

(You create the latency alias resource record sets after you create resource record sets for the individual Amazon EC2 instances.)

Within each region, you have two Amazon EC2 instances. You create a weighted resource record set for each instance. The name and the type are the same for both of the weighted resource record sets in each region.

When you have multiple resources in a region, you can create weighted or failover resource record sets for your resources. You can also create even more complex configurations by creating weighted alias or failover alias resource record sets that, in turn, refer to multiple resources.

Each weighted resource record set has an associated health check. The IP address for each health check matches the IP address for the corresponding resource record set. This isn't required, but it's the most common configuration.

For both latency alias resource record sets, you set the value of Evaluate Target Health to Yes.

You use the Evaluate Target Health setting for each latency alias resource record set to make Amazon Route 53 evaluate the health of the alias targets—the weighted resource record sets—and respond accordingly.

The preceding diagram illustrates the following sequence of events:

Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region.

Amazon Route 53 selects a weighted resource record set based on weight. Evaluate Target Health is Yes for the latency alias resource record set, so Amazon Route 53 checks the health of the selected weighted resource record set.

The health check failed, so Amazon Route 53 chooses another weighted resource record set based on weight and checks its health. That resource record set also is unhealthy.

Amazon Route 53 backs out of that branch of the tree, looks for the latency alias resource record set with the next-best latency, and chooses the resource record set for ap-southeast-2.

Amazon Route 53 again selects a resource record set based on weight, and then checks the health of the selected resource record set. The health check passed, so Amazon Route 53 returns the applicable value in response to the query.

What Happens When You Associate a Health Check with an Alias Resource Record Set?

You can associate a health check with an alias resource record set instead of or in addition to setting the value of Evaluate Target Health to Yes. However, it's generally more useful if Amazon Route 53 responds to queries based on the health of the underlying resources—the HTTP servers, database servers, and other resources that your alias resource record sets refer to. For example, suppose the following configuration:

You assign a health check to a latency alias resource record set for which the alias target is a group of weighted resource record sets.

You set the value of Evaluate Target Health to Yes for the latency alias resource record set.

In this configuration, both of the following must be true before Amazon Route 53 will return the applicable value for a weighted resource record set:

The health check associated with the latency alias resource record set must pass.

At least one weighted resource record set must be considered healthy, either because it's associated with a health check that passes or because it's not associated with a health check. In the latter case, Amazon Route 53 always considers the weighted resource record set healthy.

If the health check for the latency alias resource record set fails, Amazon Route 53 stops responding to queries using any of the weighted resource record sets in the alias target, even if they're all healthy. Amazon Route 53 doesn't know the status of the weighted resource record sets because it never looks past the failed

health check on the alias resource record set.

What Happens When You Omit Health Checks?

In a complex configuration, it's important to associate health checks with all of the non-alias resource record sets. Let's return to the preceding example, but assume that a health check is missing on one of the weighted resource record sets in the us-east-1 region:

Here's what happens when you omit a health check on a non-alias resource record set in this configuration:

Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region.

Amazon Route 53 looks up the alias target for the latency alias resource record set, and checks the status of the corresponding health checks. The health check for one weighted resource record set failed, so that resource record set is omitted from consideration.

The other weighted resource record set in the alias target for the us-east-1 region has no health check. The corresponding resource might or might not be healthy, but without a health check, Amazon Route 53 has no way to know. Amazon Route 53 assumes that the resource is healthy and returns the applicable value in response to the query.

What Happens When You Set Evaluate Target Health to No?

In general, you also want to set Evaluate Target Health to Yes for all of the alias resource record sets. In the following example, all of the weighted resource record sets have associated health checks, but Evaluate Target Health is set to No for the latency alias resource record set for the us-east-1 region:

Here's what happens when you set Evaluate Target Health to No for an alias resource record set in this configuration:

Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region.

Amazon Route 53 determines what the alias target is for the latency alias resource record set, and checks the corresponding health checks. They're both failing.

Because the value of Evaluate Target Health is No for the latency alias resource record set for the us-east-1 region, Amazon Route 53 must choose one resource record set in this branch instead of backing out of the branch and looking for a healthy resource record set in the ap-southeast-2 region.

#### NEW QUESTION 196

A customer has established an AWS Direct Connect connection to AWS. The link is up and routes are being advertised from the customer's end, however the customer is unable to connect from EC2 instances inside its VPC to servers residing in its datacenter.

Which of the following options provide a viable solution to remedy this situation? (Choose 2 answers)

- A. Add a route to the route table with an IPsec VPN connection as the target.
- B. Enable route propagation to the virtual pinnate gateway (VGW).
- C. Enable route propagation to the customer gateway (CGW).
- D. Modify the route table of all Instances using the 'route' command.
- E. Modify the Instances VPC subnet route table by adding a route back to the customer's on-premises environment.

**Answer:** AC

#### NEW QUESTION 200

A web company is looking to implement an external payment service into their highly available application deployed in a VPC Their application EC2 instances are behind a public facing ELB Auto scaling is used to add additional instances as traffic increases under normal load the application runs 2 instances in the Auto Scaling group but at peak it can scale 3x in size. The application instances need to communicate with the payment service over the Internet which requires whitelisting of all public IP addresses used to communicate with it. A maximum of 4 whitelisting IP addresses are allowed at a time and can be added through an API.

How should they architect their solution?

- A. Route payment requests through two NAT instances setup for High Availability and whitelist the Elastic IP addresses attached to the MAT instances.
- B. Whitelist the VPC Internet Gateway Public IP and route payment requests through the Internet Gateway.
- C. Whitelist the ELB IP addresses and route payment requests from the Application servers through the ELB.
- D. Automatically assign public IP addresses to the application instances in the Auto Scaling group and run a script on boot that adds each instances public IP address to the payment validation whitelist API.

**Answer:** D

#### NEW QUESTION 201

You are implementing AWS Direct Connect. You intend to use AWS public service end points such as Amazon S3, across the AWS Direct Connect link. You want other Internet traffic to use your existing link to an Internet Service Provider.

What is the correct way to configure AWS Direct connect for access to services such as Amazon S3?

- A. Configure a public Interface on your AWS Direct Connect link Configure a static route via your AWS Direct Connect link that points to Amazon S3 Advertise a default route to AWS using BGP.
- B. Create a private interface on your AWS Direct Connect link
- C. Configure a static route via your AWS Direct connect link that points to Amazon S3 Configure specific routes to your network in your VPC,
- D. Create a public interface on your AWS Direct Connect link Redistribute BGP routes into your existing routing infrastructure advertise specific routes for your network to AWS.
- E. Create a private interface on your AWS Direct connect link
- F. Redistribute BGP routes into your existing routing infrastructure and advertise a default route to AWS.

**Answer:** C

#### NEW QUESTION 202

You've been brought in as solutions architect to assist an enterprise customer with their migration of an e-commerce platform to Amazon Virtual Private Cloud (VPC) The previous architect has already deployed a 3-tier VPC, The configuration is as follows:

VPC: vpc-2f8bc447 IGW: igw-2d8bc445 NACL: ad-208bc448

Subnets and Route Tables: Web servers: subnet-258bc44d

Application servers: subnet-248bc44c Database servers: subnet-9189c6f9 Route Tables:

rtb-218bc449 rtb-238bc44b Associations:

subnet-258bc44d : rtb-218bc449 subnet-248bc44c : rtb-238bc44b subnet-9189c6f9 : rtb-238bc44b

You are now ready to begin deploying EC2 instances into the VPC Web servers must have direct access to the internet Application and database servers cannot have direct access to the internet.

Which configuration below will allow you the ability to remotely administer your application and database servers, as well as allow these servers to retrieve updates from the Internet?

- A. Create a bastion and NAT instance in subnet-258bc44d, and add a route from rtb- 238bc44b to the NAT instance.
- B. Add a route from rtb-238bc44b to igw-2d8bc445 and add a bastion and NAT instance within subnet-248bc44c.
- C. Create a bastion and NAT instance in subnet-248bc44c, and add a route from rtb- 238bc44b to subneb258bc44d.
- D. Create a bastion and NAT instance in subnet-258bc44d, add a route from rtb-238bc44b to lgw- 2d8bc445, and a new NACL that allows access between subnet-258bc44d and subnet -248bc44c.

**Answer:** A

#### NEW QUESTION 206

You're running an application on-premises due to its dependency on non-x86 hardware and want to use AWS for data backup. Your backup application is only able to write to POSIX-compatible blockbased storage. You have 140TB of data and would like to mount it as a single folder on your file server. Users must be able to access portions of this data while the backups are taking place. What backup solution would be most appropriate for this use case?

- A. Use Storage Gateway and configure it to use Gateway Cached volumes.
- B. Configure your backup software to use S3 as the target for your data backups.
- C. Configure your backup software to use Glacier as the target for your data backups.
- D. Use Storage Gateway and configure it to use Gateway Stored volume

**Answer:** A

#### Explanation:

##### Gateway-Cached Volume Architecture

Gateway-cached volumes let you use Amazon Simple Storage Service (Amazon S3) as your primary data storage while retaining frequently accessed data locally in your storage gateway. Gateway cached volumes minimize the need to scale your on-premises storage infrastructure, while still providing your applications with low-latency access to their frequently accessed data. You can create storage volumes up to 32 TiB in size and attach to them as iSCSI devices from your on-premises application servers. Your gateway stores data that you write to these volumes in Amazon S3 and retains recently read data in your on-premises storage gateway's cache and upload buffer storage.

Gateway-cached volumes can range from 1 GiB to 32 TiB in size and must be rounded to the nearest GiB. Each gateway configured for gateway-cached volumes can support up to 32 volumes for a total maximum storage volume of 1,024 TiB (1 PiB).

In the gateway-cached volume solution, AWS Storage Gateway stores all your on-premises application data in a storage volume in Amazon S3.

The following diagram provides an overview of the AWS Storage Gateway-cached volume deployment.

After you've installed the AWS Storage Gateway software appliance-the virtual machine (VM)-on a host in your data center and activated it, you can use the AWS Management Console to provision storage

volumes backed by Amazon S3. You can also provision storage volumes programmatically using the AWS Storage Gateway API or the AWS SDK libraries. You then mount these storage volumes to your on-premises application servers as iSCSI devices.

You also allocate disks on-premises for the VM. These on-premises disks serve the following purposes: Disks for use by the gateway as cache storage - As your applications write data to the storage volumes in AWS, the gateway initially stores the data on the on-premises disks referred to as cache storage before uploading the data to Amazon S3. The cache storage acts as the on-premises durable store for data that is waiting to upload to Amazon S3 from the upload buffer.

The cache storage also lets the gateway store your application's recently accessed data on-premises for low-latency access. If your application requests data, the gateway first checks the cache storage for the data before checking Amazon S3.

You can use the following guidelines to determine the amount of disk space to allocate for cache storage. Generally, you should allocate at least 20 percent of your existing file store size as cache storage. Cache storage should also be larger than the upload buffer. This latter guideline helps ensure cache storage is large enough to persistently hold all data in the upload buffer that has not yet been uploaded to Amazon S3.

Disks for use by the gateway as the upload buffer - To prepare for upload to Amazon S3, your gateway also stores incoming data in a staging area, referred to as an upload buffer. Your gateway uploads this buffer data over an encrypted Secure Sockets Layer (SSL) connection to AWS, where it is stored encrypted in Amazon S3.

You can take incremental backups, called snapshots, of your storage volumes in Amazon S3. These point-in-time snapshots are also stored in Amazon S3 as Amazon EBS snapshots. When you take a new snapshot, only the data that has changed since your last snapshot is stored. You can initiate snapshots on a scheduled or one-time basis. When you delete a snapshot, only the data not needed for any other snapshots is removed.

You can restore an Amazon EBS snapshot to a gateway storage volume if you need to recover a backup of your data. Alternatively, for snapshots up to 16 TiB in size, you can use the snapshot as a starting point for a new Amazon EBS volume. You can then attach this new Amazon EBS volume to an Amazon EC2 instance.

All gateway-cached volume data and snapshot data is stored in Amazon S3 encrypted at rest using server-side encryption (SSE). However, you cannot access this data with the Amazon S3 API or other tools such as the Amazon S3 console.

#### NEW QUESTION 208

Your department creates regular analytics reports from your company's log files. All log data is collected in Amazon S3 and processed by daily Amazon Elastic MapReduce (EMR) jobs that generate daily PDF reports and aggregated tables in CSV format for an Amazon Redshift data warehouse.

Your CFO requests that you optimize the cost structure for this system.

Which of the following alternatives will lower costs without compromising average performance of the system or data integrity for the raw data?

- A. Use reduced redundancy storage (RRS) for all data in S3. Use a combination of Spot Instances and Reserved Instances for Amazon EMR job
- B. Use Reserved Instances for Amazon Redshift.
- C. Use reduced redundancy storage (RRS) for PDF and .csv data in S3. Add Spot Instances to EMR job
- D. Use Spot Instances for Amazon Redshift.
- E. Use reduced redundancy storage (RRS) for PDF and .csv data in Amazon S3. Add Spot Instances to Amazon EMR job
- F. Use Reserved Instances for Amazon Redshift.
- G. Use reduced redundancy storage (RRS) for all data in Amazon S3. Add Spot Instances to Amazon EMR job
- H. Use Reserved Instances for Amazon Redshift.

**Answer:** C

#### Explanation:

##### Using Reduced Redundancy Storage

Amazon S3 stores objects according to their storage class. It assigns the storage class to an object when it is written to Amazon S3. You can assign objects a specific storage class (standard or reduced redundancy) only when you write the objects to an Amazon S3 bucket or when you copy objects that are already stored in Amazon S3. Standard is the default storage class. For information about storage classes, see Object Key and Metadata.

In order to reduce storage costs, you can use reduced redundancy storage for noncritical, reproducible data at lower levels of redundancy than Amazon S3 provides with standard storage. The lower level of redundancy results in less durability and availability, but in many cases, the lower costs can make reduced redundancy storage an acceptable storage solution. For example, it can be a cost effective solution for sharing media content that is durably stored elsewhere. It can also make sense if you are storing thumbnails and other resized images that can be easily reproduced from an original image. Reduced

redundancy storage is designed to provide 99.99% durability of objects over a given year.

This durability level corresponds to an average annual expected loss of 0.01% of objects. For example, if you store 10,000 objects using the RRS option, you can, on average, expect to incur an annual loss of a single object per year (0.01% of 10,000 objects).

Note

This annual loss represents an expected average and does not guarantee the loss of less than 0.01% of objects in a given year.

Reduced redundancy storage stores objects on multiple devices across multiple facilities, providing 400 times the durability of a typical disk drive, but it does not replicate objects as many times as Amazon S3 standard storage. In addition, reduced redundancy storage is designed to sustain the loss of data in a single facility. If an object in reduced redundancy storage has been lost, Amazon S3 will return a 405 error on requests made to that object. Amazon S3 also offers notifications for reduced redundancy storage object loss: you can configure your bucket so that when Amazon S3 detects the loss of an RRS object, a notification will be sent through Amazon Simple Notification Service (Amazon SNS). You can then replace the lost object. To enable notifications, you can use the Amazon S3 console to set the Notifications property of your bucket.

#### NEW QUESTION 212

You are the new IT architect in a company that operates a mobile sleep tracking application

When activated at night, the mobile app is sending collected data points of 1 kilobyte every 5 minutes to your backend

The backend takes care of authenticating the user and writing the data points into an Amazon DynamoDB table.

Every morning, you scan the table to extract and aggregate last night's data on a per user basis, and store the results in Amazon S3.

Users are notified via Amazon SNS mobile push notifications that new data is available, which is parsed and visualized by (The mobile app Currently you have around 100k users who are mostly based out of North America.

You have been tasked to optimize the architecture of the backend system to lower cost what would you recommend? (Choose 2 answers)

- A. Create a new Amazon DynamoDB (able each day and drop the one for the previous day after its data is on Amazon S3.
- B. Have the mobile app access Amazon DynamoDB directly instead of JSON files stored on Amazon S3.
- C. Introduce an Amazon SQS queue to buffer writes to the Amazon DynamoDB table and reduce provisioned write throughput.
- D. Introduce Amazon ElastiCache to cache reads from the Amazon DynamoDB table and reduce provisioned read throughput.
- E. Write data directly into an Amazon Redshift cluster replacing both Amazon DynamoDB and Amazon S3.

**Answer: BD**

#### NEW QUESTION 217

Your website is serving on-demand training videos to your workforce. Videos are uploaded monthly in high resolution MP4 format. Your workforce is distributed globally often on the move and using company-provided tablets that require the HTTP Live Streaming (HLS) protocol to watch a video. Your company has no video transcoding expertise and it required you may need to pay for a consultant.

How do you implement the most cost-efficient architecture without compromising high availability and quality of video delivery?

- A. A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queue
- B. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few days
- C. CloudFront to serve HLS transcoded videos from EC2.
- D. Elastic Transcoder to transcode original high-resolution MP4 videos to HL
- E. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few days
- F. CloudFront to serve HLS transcoded videos from EC2.
- G. Elastic Transcoder to transcode original high-resolution MP4 videos to HL
- H. S3 to host videos with Lifecycle Management to archive original files to Glacier after a few days
- I. CloudFront to serve HLS transcoded videos from S3.
- J. A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queue
- K. S3 to host videos with Lifecycle Management to archive all files to Glacier after a few days
- L. CloudFront to serve HLS transcoded videos from Glacier.

**Answer: C**

#### NEW QUESTION 218

You currently operate a web application in the AWS US-East region. The application runs on an autoscaled layer of EC2 instances and an RDS Multi-AZ database. Your IT security compliance officer has tasked you to develop a reliable and durable logging solution to track changes made to your EC2, IAM and RDS resources. The solution must ensure the integrity and confidentiality of your log data. Which of these solutions would you recommend?

- A. Create a new CloudTrail trail with one new S3 bucket to store the logs and with the global services option selected. Use IAM roles, S3 bucket policies and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs.
- B. Create a new CloudTrail with one new S3 bucket to store the logs. Configure SNS to send log file delivery notifications to your management system. Use IAM roles and S3 bucket policies on the S3 bucket that stores your logs.
- C. Create a new CloudTrail trail with an existing S3 bucket to store the logs and with the global services option selected. Use S3 ACLs and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs.
- D. Create three new CloudTrail trails with three new S3 buckets to store the logs: one for the AWS Management console, one for AWS SDKs and one for command line tools. Use IAM roles and S3 bucket policies on the S3 buckets that store your logs.

**Answer: A**

#### NEW QUESTION 222

You are designing a data leak prevention solution for your VPC environment. You want your VPC instances to be able to access software depots and distributions on the Internet for product updates. The depots and distributions are accessible via third party CDNs by their URLs. You want to explicitly deny any other outbound connections from your VPC instances to hosts on the Internet.

Which of the following options would you consider?

- A. Configure a web proxy server in your VPC and enforce URL-based rules for outbound access. Remove default routes.
- B. Implement security groups and configure outbound rules to only permit traffic to software depots.
- C. Move all your instances into private VPC subnets, remove default routes from all routing tables and add specific routes to the software depots and distributions only.
- D. Implement network access control lists to all specific destinations, with an implicit deny as a rule.

**Answer:** A

**NEW QUESTION 224**

An administrator is using Amazon CloudFormation to deploy a three tier web application that consists of a web tier and application tier that will utilize Amazon DynamoDB for storage when creating the CloudFormation template which of the following would allow the application instance access to the DynamoDB tables without exposing API credentials?

- A. Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and associate the Role to the application instances by referencing an instance profile.
- B. Use the Parameter section in the Cloud Formation template to have the user input Access and Secret Keys from an already created IAM user that has the permissions required to read and write from the required DynamoDB table.
- C. Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and reference the Role in the instance profile property of the application instance.
- D. Create an identity and Access Management user in the CloudFormation template that has permissions to read and write from the required DynamoDB table, use the GetAtt function to retrieve the Access and secret keys and pass them to the application instance through user-data.

**Answer:** C

**NEW QUESTION 227**

You have an application running on an EC2 Instance which will allow users to download files from a private S3 bucket using a pre-assigned URL. Before generating the URL the application should verify the existence of the file in S3. How should the application use AWS credentials to access the S3 bucket securely?

- A. Use the AWS account access Keys the application retrieves the credentials from the source code of the application.
- B. Create an IAM user for the application with permissions that allow list access to the S3 bucket launch the instance as the IAM user and retrieve the IAM user's credentials from the EC2 instance user data.
- C. Create an IAM role for EC2 that allows list access to objects in the S3 bucket
- D. Launch the instance with the role, and retrieve the role's credentials from the EC2 Instance metadata
- E. Create an IAM user for the application with permissions that allow list access to the S3 bucket
- F. The application retrieves the IAM user credentials from a temporary directory with permissions that allow read access only to the application user.

**Answer:** C

**NEW QUESTION 229**

A benefits enrollment company is hosting a 3-tier web application running in a VPC on AWS which includes a NAT (Network Address Translation) instance in the public Web tier. There is enough provisioned capacity for the expected workload for the new fiscal year benefit enrollment period plus some extra overhead. Enrollment proceeds nicely for two days and then the web tier becomes unresponsive, upon investigation using CloudWatch and other monitoring tools it is discovered that there is an extremely large and unanticipated amount of inbound traffic coming from a set of 15 specific IP addresses over port 80 from a country where the benefits company has no customers. The web tier instances are so overloaded that benefit enrollment administrators cannot even SSH into them. Which actMty would be useful in defending against this attack?

- A. Create a custom route table associated with the web tier and block the attacking IP addresses from the IGW (Internet Gateway)
- B. Change the EIP (Elastic IP Address) of the NAT instance in the web tier subnet and update the Main Route Table with the new EIP
- C. Create 15 Security Group rules to block the attacking IP addresses over port 80
- D. Create an inbound NACL (Network Access control list) associated with the web tier subnet with deny rules to block the attacking IP addresses

**Answer:** D

**Explanation:**

Use AWS Identity and Access Management (IAM) to control who in your organization has permission to create and manage security groups and network ACLs (NACL). Isolate the responsibilities and roles for better defense. For example, you can give only your network administrators or security admin the permission to manage the security groups and restrict other roles.

**NEW QUESTION 234**

Your fortune 500 company has undertaken a TCO analysis evaluating the use of Amazon S3 versus acquiring more hardware. The outcome was that all employees would be granted access to use Amazon S3 for storage of their personal documents. Which of the following will you need to consider so you can set up a solution that incorporates single sign-on from your corporate AD or LDAP directory and restricts access for each user to a designated user folder in a bucket? (Choose 3 Answers)

- A. Setting up a federation proxy or identity provider
- B. Using AWS Security Token Service to generate temporary tokens
- C. Tagging each folder in the bucket
- D. Configuring IAM role
- E. Setting up a matching IAM user for every user in your corporate directory that needs access to a folder in the bucket

**Answer:** ABD

**NEW QUESTION 237**

Your company policies require encryption of sensitive data at rest. You are considering the possible options for protecting data while storing it at rest on an EBS data volume, attached to an EC2 instance. Which of these options would allow you to encrypt your data at rest? (Choose 3 answers)

- A. Implement third party volume encryption tools
- B. Do nothing as EBS volumes are encrypted by default
- C. Encrypt data inside your applications before storing it on EBS
- D. Encrypt data using native data encryption drivers at the file system level
- E. Implement SSL/TLS for all services running on the server

**Answer:** ACD

#### NEW QUESTION 238

You have a periodic image analysis application that gets some files in input analyzes them and for each file writes some data in output to a text file the number of files in input per day is high and concentrated in a few hours of the day.

Currently you have a server on EC2 with a large EBS volume that hosts the input data and the results it takes almost 20 hours per day to complete the process. What services could be used to reduce the elaboration time and improve the availability of the solution?

- A. S3 to store I/O file
- B. SQS to distribute elaboration commands to a group of hosts working in parallel
- C. Auto scaling to dynamically size the group of hosts depending on the length of the SQS queue
- D. EBS with Provisioned IOPS (PIOPS) to store I/O file
- E. SNS to distribute elaboration commands to a group of hosts working in parallel Auto Scaling to dynamically size the group of hosts depending on the number of SNS notifications
- F. S3 to store I/O files, SNS to distribute elaboration commands to a group of hosts working in parallel
- G. Auto scaling to dynamically size the group of hosts depending on the number of SNS notifications
- H. EBS with Provisioned IOPS (PIOPS) to store I/O files SQS to distribute elaboration commands to a group of hosts working in parallel Auto Scaling to dynamically size the group of hosts depending on the length of the SQS queue.

**Answer:** D

#### Explanation:

Amazon EBS allows you to create storage volumes and attach them to Amazon EC2 instances. Once attached, you can create a file system on top of these volumes, run a database, or use them in any other way you would use a block device. Amazon EBS volumes are placed in a specific Availability Zone, where they are automatically replicated to protect you from the failure of a single component.

Amazon EBS provides three volume types: General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic. The three volume types differ in performance characteristics and cost, so you can choose the right storage performance and price for the needs of your applications. All EBS volume types offer the same durable snapshot capabilities and are designed for 99.999% availability.

#### NEW QUESTION 241

A large real-estate brokerage is exploring the option of adding a cost-effective location-based alert to their existing mobile application. The application backend infrastructure currently runs on AWS. Users who opt in to this service will receive alerts on their mobile device regarding real-estate offers in proximity to their location. For the alerts to be relevant, delivery time needs to be in the low minute count. The existing mobile app has 5 million users across the US. Which one of the following architectural suggestions would you make to the customer?

- A. The mobile application will submit its location to a web service endpoint utilizing Elastic Load Balancing and EC2 instances. DynamoDB will be used to store and retrieve relevant offers. EC2 instances will communicate with mobile carriers/device providers to push alerts back to mobile application.
- B. Use AWS DirectConnect or VPN to establish connectivity with mobile carriers. EC2 instances will receive the mobile applications' location through carrier connection. ROS will be used to store and retrieve relevant offers. EC2 instances will communicate with mobile carriers to push alerts back to the mobile application.
- C. The mobile application will send device location using SNS.
- D. EC2 instances will retrieve the relevant offers from DynamoDB. AWS Mobile Push will be used to send offers to the mobile application.
- E. The mobile application will send device location using AWS Mobile Push. EC2 instances will retrieve the relevant offers from DynamoDB. EC2 instances will communicate with mobile carriers/device providers to push alerts back to the mobile application.

**Answer:** A

#### NEW QUESTION 243

A company is building a voting system for a popular TV show. Viewers watch the performances then visit the show's website to vote for their favorite performer. It is expected that in a short period of time after the show has finished the site will receive millions of visitors. The visitors will first login to the site using their Amazon.com credentials and then submit their vote. After the voting is completed the page will display the vote totals. The company needs to build the site such that it can handle the rapid influx of traffic while maintaining good performance but also wants to keep costs to a minimum. Which of the design patterns below should they use?

- A. Use CloudFront and an Elastic Load balancer in front of an auto-scaled set of web servers, the web servers will first call the Login With Amazon service to authenticate the user then process the user's vote and store the result into a multi-AZ Relational Database Service instance.
- B. Use CloudFront and the static website hosting feature of S3 with the Javascript SDK to call the Login With Amazon service to authenticate the user, use IAM Roles to gain permissions to a DynamoDB table to store the user's vote.
- C. Use CloudFront and an Elastic Load Balancer in front of an auto-scaled set of web servers, the web servers will first call the Login with Amazon service to authenticate the user, the web servers will process the user's vote and store the result into a DynamoDB table using IAM Roles for EC2 instances to gain permissions to the DynamoDB table.
- D. Use CloudFront and an Elastic Load Balancer in front of an auto-scaled set of web servers, the web servers will first call the Login With Amazon service to authenticate the user, the web servers will process the user's vote and store the result into an SQS queue using IAM Roles for EC2 instances to gain permissions to the SQS queue.
- E. A set of application servers will then retrieve the items from the queue and store the result into a DynamoDB table.

**Answer:** D

#### NEW QUESTION 248

Your team has a Tomcat-based Java application you need to deploy into development, test, and production environments. After some research, you opt to use Elastic Beanstalk due to its tight integration with your developer tools and RDS due to its ease of management. Your QA team lead points out that you need to roll a sanitized set of production data into your environment on a nightly basis. Similarly, other software teams in your org want access to that same restored data via their EC2 instances in your VPC. The

optimal setup for persistence and security that meets the above requirements would be the following.

- A. Create your RDS instance as part of your Elastic Beanstalk definition and alter its security group to allow access to it from hosts in your application subnets.
- B. Create your RDS instance separately and add its IP address to your application's DB connection strings in your code. Alter its security group to allow access to it from hosts within your VPC's IP address block.
- C. Create your RDS instance separately and pass its DNS name to your app's DB connection string as an environment variable.

- D. Create a security group for client machines and add it as a valid source for DB traffic to the security group of the RDS instance itself.
- E. Create your RDS instance separately and pass its DNS name to your's DB connection string as an environment variable Alter its security group to allow access to It from hosts In your application subnets.

**Answer:** A

**NEW QUESTION 250**

What does Amazon S3 stand for?

- A. Simple Storage Solution.
- B. Storage Storage Storage (triple redundancy Storage).
- C. Storage Sewer Solution.
- D. Simple Storage Sewic

**Answer:** D

**NEW QUESTION 251**

Before I delete an EBS volume, what can I do if I want to recreate the volume later?

- A. Create a copy of the EBS volume (not a snapshot)
- B. Store a snapshot of the volume
- C. Download the content to an EC2 instance
- D. Back up the data in to a physical disk

**Answer:** B

**NEW QUESTION 256**

Select the most correct

The device name /dev/sdal (within Amazon EC2) is \_

- A. Possible for EBS volumes
- B. Reserved for the root device
- C. Recommended for EBS volumes
- D. Recommended for instance store volumes

**Answer:** B

**NEW QUESTION 261**

If I want an instance to have a public IP address, which IP address should I use'?

- A. Elastic I P Address
- B. Class B IP Address
- C. Class A IP Address
- D. Dynamic IP Address

**Answer:** A

**NEW QUESTION 266**

What does RRS stand for when talking about S3?

- A. Redundancy Removal System
- B. Relational Rights Storage
- C. Regional Rights Standard
- D. Reduced Redundancy Storage

**Answer:** D

**NEW QUESTION 268**

Fill in the blanks: Resources that are created in AWS are identified by a unique identifier called an

- A. Amazon Resource Number
- B. Amazon Resource Nametag
- C. Amazon Resource Name
- D. Amazon Resource Namespace

**Answer:** C

**NEW QUESTION 272**

Can you create IAM security credentials for existing users?

- A. Yes, existing users can have security credentials associated with their account.
- B. No, IAM requires that all users who have credentials set up are not existing users
- C. No, security credentials are created within GROUPS, and then users are associated to GROUPS at a later time.
- D. Yes, but only IAM credentials, not ordinary security credential

**Answer:** A

**NEW QUESTION 274**

Amazon SWF is designed to help users

- A. Design graphical user interface interactions
- B. Manage user identification and authorization
- C. Store Web content
- D. Coordinate synchronous and asynchronous tasks which are distributed and fault tolerant

**Answer:** D

**NEW QUESTION 277**

What are the initial settings of a user created security group?

- A. Allow all inbound traffic and Allow no outbound traffic
- B. Allow no inbound traffic and Allow no outbound traffic
- C. Allow no inbound traffic and Allow all outbound traffic
- D. Allow all inbound traffic and Allow all outbound traffic

**Answer:** C

**NEW QUESTION 278**

Will my standby RDS instance be in the same Region as my primary?

- A. Only for Oracle RDS types
- B. Yes
- C. Only if configured at launch
- D. No

**Answer:** B

**NEW QUESTION 280**

Typically, you want your application to check whether a request generated an error before you spend any time processing results. The easiest way to find out if an error occurred is to look for an `_` node in the response from the Amazon RDS API.

- A. Incorrect
- B. Error
- C. FALSE

**Answer:** B

**NEW QUESTION 282**

You must increase storage size in increments of at least `_` %

- A. 40
- B. 20
- C. 50
- D. 10

**Answer:** D

**NEW QUESTION 285**

Using Amazon CloudWatch's Free Tier, what is the frequency of metric updates which you receive?

- A. 5 minutes
- B. 500 milliseconds.
- C. 30 seconds
- D. 1 minute

**Answer:** A

**NEW QUESTION 289**

What are the Amazon EC2 API tools?

- A. They don't exist
- B. The Amazon EC2 CLI tools, instead, are used to manage permissions.
- C. Command-line tools to the Amazon EC2 web service.
- D. They are a set of graphical tools to manage EC2 instances.
- E. They don't exist
- F. The Amazon API tools are a client interface to Amazon Web Services

**Answer:** B

**NEW QUESTION 294**

EBS Snapshots occur \_

- A. Asynchronously
- B. Synchronously
- C. Weekly

**Answer:** A

**NEW QUESTION 297**

Which of the following cannot be used in Amazon EC2 to control who has access to specific Amazon EC2 instances?

- A. Security Groups
- B. IAM System
- C. SSH keys
- D. Windows passwords

**Answer:** B

**NEW QUESTION 300**

Can a 'user' be associated with multiple AWS accounts?

- A. No
- B. Yes

**Answer:** A

**NEW QUESTION 305**

True or False: Manually created DB Snapshots are deleted after the DB Instance is deleted.

- A. TRUE
- B. FALSE

**Answer:** A

**NEW QUESTION 306**

What is Amazon Glacier?

- A. You mean Amazon "Iceberg": it's a low-cost storage service.
- B. A security tool that allows to "freeze" an EBS volume and perform computer forensics on it.
- C. A low-cost storage service that provides secure and durable storage for data archMng and backup.
- D. It's a security tool that allows to "freeze" an EC2 instance and perform computer forensics on i

**Answer:** C

**NEW QUESTION 311**

What is the durability of 53 RRS?

- A. 99.99%
- B. 99.95%
- C. 99.995%
- D. 99.999999999%

**Answer:** A

**NEW QUESTION 316**

While launching an RDS DB instance, on which page I can select the Availability Zone?

- A. REVIEW
- B. DB INSTANCE DETAILS
- C. MANAGEMENT OPTIONS
- D. ADDITIONAL CONFIGURATION

**Answer:** D

**NEW QUESTION 317**

IAM's Policy Evaluation Logic always starts with a default \_ for every request, except for those that use the AWS account's root security credentials b

- A. Permit
- B. Deny
- C. Cancel

**Answer:** B

**NEW QUESTION 322**

By default, when an EBS volume is attached to a Windows instance, it may show up as any drive letter on the instance. You can change the settings of the \_ Senrice to set the drive letters of the EBS volumes per your specifications.

- A. EBS Config Senrice
- B. AMI Config Senrice
- C. Ec2 Config Senrice
- D. Ec2-AMI Config Senrice

**Answer: C**

**NEW QUESTION 323**

For each DB Instance class, what is the maximum size of associated storage capacity?

- A. 5GB
- B. 1 TB
- C. 2TB
- D. 500GB

**Answer: B**

**NEW QUESTION 324**

To view information about an Amazon EBS volume, open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>, click in the Navigation pane.

- A. EBS
- B. Describe
- C. Details
- D. Volumes

**Answer: D**

**NEW QUESTION 329**

While creating the snapshots using the API, which Action should I be using?

- A. MakeSnapShot
- B. FreshSnapshot
- C. DeploySnapshot
- D. CreateSnapshot

**Answer: D**

**NEW QUESTION 330**

What happens to the I/O operations while you take a database snapshot?

- A. I/O operations to the database are suspended for a few minutes while the backup is in progress.
- B. I/O operations to the database are sent to a Replica (if available) for a few minutes while the backup is in progress.
- C. I/O operations will be functioning normally
- D. I/O operations to the database are suspended for an hour while the backup is in progress

**Answer: A**

**NEW QUESTION 332**

Read Replicas require a transactional storage engine and are only supported for the \_\_ storage engine

- A. OracleI<sup>S</sup>AM
- B. MSSQLDB
- C. InnoDB
- D. IV|y|SAIV|

**Answer: C**

**NEW QUESTION 337**

Which service enables AWS customers to manage users and permissions in AWS?

- A. AWS Access Control Service (ACS)
- B. AWS Identity and Access Management (IAM)
- C. AWS Identity Manager (AIM)

**Answer: B**

**NEW QUESTION 342**

Which Amazon Storage behaves like raw, unformatted, external block devices that you can attach to your instances?

- A. None of these.
- B. Amazon Instance Storage

- C. Amazon EBS
- D. All of these

**Answer:** C

**NEW QUESTION 344**

MySQL installations default to port \_\_.

- A. A.3306B.443
- B. 80
- C. 1158

**Answer:** A

**NEW QUESTION 349**

If you have chosen Multi-AZ deployment, in the event of a planned or unplanned outage of your primary DB Instance, Amazon RDS automatically switches to the standby replica. The automatic failover mechanism simply changes the record of the main DB Instance to point to the standby DB Instance.

- A. DNAME
- B. CNAME
- C. TXT
- D. MX

**Answer:** B

**NEW QUESTION 351**

If I modify a DB Instance or the DB parameter group associated with the instance, should I reboot the instance for the changes to take effect?

- A. No
- B. Yes

**Answer:** B

**NEW QUESTION 353**

Can I test my DB Instance against a new version before upgrading?

- A. No
- B. Yes
- C. Only in VPC

**Answer:** B

**NEW QUESTION 354**

Can I encrypt connections between my application and my DB Instance using SSL?

- A. No
- B. Yes
- C. Only in VPC
- D. Only in certain regions

**Answer:** B

**NEW QUESTION 356**

What are the four levels of AWS Premium Support?

- A. Basic, Developer, Business, Enterprise
- B. Basic, Startup, Business, Enterprise
- C. Free, Bronze, Silver, Gold
- D. All support is free

**Answer:** A

**NEW QUESTION 361**

Can the string value of 'Key' be prefixed with :aws:"?

- A. Only in GovCloud
- B. Only for 53 not EC2
- C. Yes
- D. No

**Answer:** D

**NEW QUESTION 362**

The new DB Instance that is created when you promote a Read Replica retains the backup window period.

- A. TRUE
- B. FALSE

**Answer:** A

**NEW QUESTION 363**

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

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

\* One year free update

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

\* Trusted by Millions

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

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your AWS-Solution-Architect-Associate Exam with Our Prep Materials Via below:**

<https://www.certleader.com/AWS-Solution-Architect-Associate-dumps.html>