

## 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**

You have set up an Auto Scaling group. The cool down period for the Auto Scaling group is 7 minutes. The first instance is launched after 3 minutes, while the second instance is launched after 4 minutes. How many minutes after the first instance is launched will Auto Scaling accept another scaling activity request?

- A. 11 minutes
- B. 7 minutes
- C. 10 minutes
- D. 14 minutes

**Answer: A**

**Explanation:**

If an Auto Scaling group is launching more than one instance, the cool down period for each instance starts after that instance is launched. The group remains locked until the last instance that was launched has completed its cool down period. In this case the cool down period for the first instance starts after 3 minutes and finishes at the 10th minute (3+7 cool down), while for the second instance it starts at the 4th minute and finishes at the 11th minute (4+7 cool down). Thus, the Auto Scaling group will receive another request only after 11 minutes.

Reference: [http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS\\_Concepts.html](http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html)

**NEW QUESTION 3**

In Amazon EC2 Container Service components, what is the name of a logical grouping of container instances on which you can place tasks?

- A. A cluster
- B. A container instance
- C. A container
- D. A task definition

**Answer: A**

**Explanation:**

Amazon ECS contains the following components:

A Cluster is a logical grouping of container instances that you can place tasks on.

A Container instance is an Amazon EC2 instance that is running the Amazon ECS agent and has been registered into a cluster.

A Task definition is a description of an application that contains one or more container definitions. A Scheduler is the method used for placing tasks on container instances.

A Service is an Amazon ECS service that allows you to run and maintain a specified number of instances of a task definition simultaneously.

A Task is an instantiation of a task definition that is running on a container instance. A Container is a Linux container that was created as part of a task.

Reference: <http://docs.aws.amazon.com/AmazonECS/latest/developerguide/Welcome.html>

**NEW QUESTION 4**

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 5**

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

- A. Yes, if configured with the Launch Config

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

**Answer:** C

**Explanation:**

The user can get notifications using SNS if he has configured the notifications while creating the Auto Scaling group.  
Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/GettingStartedTutorial.html>

**NEW QUESTION 6**

Amazon EBS provides the ability to create backups of any Amazon EC2 volume into what is known as

- A. snapshots
- B. images
- C. instance backups
- D. mirrors

**Answer:** A

**Explanation:**

Amazon allows you to make backups of the data stored in your EBS volumes through snapshots that can later be used to create a new EBS volume.  
Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Storage.html>

**NEW QUESTION 7**

After you recommend Amazon Redshift to a client as an alternative solution to paying data warehouses to analyze his data, your client asks you to explain why you are recommending Redshift. Which of the following would be a reasonable response to his request?

- A. It has high performance at scale as data and query complexity grows.
- B. It prevents reporting and analytic processing from interfering with the performance of OLTP workloads.
- C. You don't have the administrative burden of running your own data warehouse and dealing with setup, durability, monitoring, scaling, and patching.
- D. All answers listed are a reasonable response to his QUESTION

**Answer:** D

**Explanation:**

Amazon Redshift delivers fast query performance by using columnar storage technology to improve I/O efficiency and parallelizing queries across multiple nodes. Redshift uses standard PostgreSQL JDBC and ODBC drivers, allowing you to use a wide range of familiar SQL clients. Data load speed scales linearly with cluster size, with integrations to Amazon S3, Amazon DynamoDB, Amazon Elastic MapReduce, Amazon Kinesis or any SSH-enabled host.  
AWS recommends Amazon Redshift for customers who have a combination of needs, such as: High performance at scale as data and query complexity grows  
Desire to prevent reporting and analytic processing from interfering with the performance of OLTP workloads  
Large volumes of structured data to persist and query using standard SQL and existing BI tools  
Desire to the administrative burden of running one's own data warehouse and dealing with setup, durability, monitoring, scaling and patching  
Reference: [https://aws.amazon.com/running\\_databases/#redshift\\_anchor](https://aws.amazon.com/running_databases/#redshift_anchor)

**NEW QUESTION 8**

One of the criteria for a new deployment is that the customer wants to use AWS Storage Gateway. However you are not sure whether you should use gateway-cached volumes or gateway-stored volumes or even what the differences are. Which statement below best describes those differences?

- A. Gateway-cached lets you store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally
- B. Gateway-stored enables you to configure your on-premises gateway to store all your data locally and then asynchronously back up point-in-time snapshots of this data to Amazon S3.
- C. Gateway-cached is free whilst gateway-stored is not.
- D. Gateway-cached is up to 10 times faster than gateway-stored.
- E. Gateway-stored lets you store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally
- F. Gateway-cached enables you to configure your on-premises gateway to store all your data locally and then asynchronously back up point-in-time snapshots of this data to Amazon S3.

**Answer:** A

**Explanation:**

Volume gateways provide cloud-backed storage volumes that you can mount as Internet Small Computer System Interface (iSCSI) devices from your on-premises application servers. The gateway supports the following volume configurations:  
Gateway-cached volumes — You store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally. Gateway-cached volumes offer a substantial cost savings on primary storage and minimize the need to scale your storage on-premises. You also retain low-latency access to your frequently accessed data.  
Gateway-stored volumes — If you need low-latency access to your entire data set, you can configure your on-premises gateway to store all your data locally and then asynchronously back up point-in-time snapshots of this data to Amazon S3. This configuration provides durable and inexpensive off-site backups that you can recover to your local data center or Amazon EC2. For example, if you need replacement capacity for disaster recovery, you can recover the backups to Amazon EC2.  
Reference: <http://docs.aws.amazon.com/storagegateway/latest/userguide/volume-gateway.html>

**NEW QUESTION 9**

A user wants to use an EBS-backed Amazon EC2 instance for a temporary job. Based on the input data, the job is most likely to finish within a week. Which of the following steps should be followed to terminate the instance automatically once the job is finished?

- A. Configure the EC2 instance with a stop instance to terminate it.
- B. Configure the EC2 instance with ELB to terminate the instance when it remains idle.

- C. Configure the CloudWatch alarm on the instance that should perform the termination action once the instance is idle.
- D. Configure the Auto Scaling schedule actMty that terminates the instance after 7 day

**Answer: C**

**Explanation:**

Auto Scaling can start and stop the instance at a pre-defined time. Here, the total running time is unknown. Thus, the user has to use the CloudWatch alarm, which monitors the CPU utilization. The user can create an alarm that is triggered when the average CPU utilization percentage has been lower than 10 percent for 24 hours, signaling that it is idle and no longer in use. When the utilization is below the threshold limit, it will terminate the instance as a part of the instance action.

Reference: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/UsingAlarmActions.html>

**NEW QUESTION 10**

Which of the following is true of Amazon EC2 security group?

- A. You can modify the outbound rules for EC2-Classic.
- B. You can modify the rules for a security group only if the security group controls the traffic for just one instance.
- C. You can modify the rules for a security group only when a new instance is created.
- D. You can modify the rules for a security group at any tim

**Answer: D**

**Explanation:**

A security group acts as a virtual firewall that controls the traffic for one or more instances. When you launch an instance, you associate one or more security groups with the instance. You add rules to each security group that allow traffic to or from its associated instances. You can modify the rules for a security group at any time; the new rules 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 10**

An Elastic IP address (EIP) is a static IP address designed for dynamic cloud computing. With an EIP, you can mask the failure of an instance or software by rapidly remapping the address to another instance in your account. Your EIP is associated with your AWS account, not a particular EC2 instance, and it remains associated with your account until you choose to explicitly release it. By default how many EIPs is each AWS account limited to on a per region basis?

- A. 1
- B. 5
- C. Unlimited
- D. 10

**Answer: B**

**Explanation:**

By default, all AWS accounts are limited to 5 Elastic IP addresses per region for each AWS account, because public (IPv4) Internet addresses are a scarce public resource. AWS strongly encourages you to use an EIP primarily for load balancing use cases, and use DNS hostnames for all other inter-node communication. If you feel your architecture warrants additional EIPs, you would need to complete the Amazon EC2 Elastic IP Address Request Form and give reasons as to your need for additional addresses. Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html#using-instance-addressing-limit>

**NEW QUESTION 14**

In EC2, what happens to the data in an instance store if an instance reboots (either intentionally or unintentionally)?

- A. Data is deleted from the instance store for security reasons.
- B. Data persists in the instance store.
- C. Data is partially present in the instance store.
- D. Data in the instance store will be los

**Answer: B**

**Explanation:**

The data in an instance store persists only during the lifetime of its associated instance. If an instance reboots (intentionally or unintentionally), data in the instance store persists. However, data on instance store volumes is lost under the following circumstances.

Failure of an underlying drive

Stopping an Amazon EBS-backed instance Terminating an instance

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

**NEW QUESTION 19**

You are setting up a VPC and you need to set up a public subnet within that VPC. Which following requirement must be met for this subnet to be considered a public subnet?

- A. Subnet's traffic is not routed to an internet gateway but has its traffic routed to a virtual private gateway.
- B. Subnet's traffic is routed to an internet gateway.
- C. Subnet's traffic is not routed to an internet gateway.
- D. None of these answers can be considered a public subne

**Answer: B**

**Explanation:**

A virtual private cloud (VPC) is a virtual network dedicated to your AWS account. It is logically isolated from other virtual networks in the AWS cloud. You can launch your AWS resources, such as Amazon EC2 instances, into your VPC. You can configure your VPC: you can select its IP address range, create subnets,

and configure route tables, network gateways, and security settings.

A subnet is a range of IP addresses in your VPC. You can launch AWS resources into a subnet that you select. Use a public subnet for resources that must be connected to the internet, and a private subnet for resources that won't be connected to the Internet.

If a subnet's traffic is routed to an internet gateway, the subnet is known as a public subnet.

If a subnet doesn't have a route to the internet gateway, the subnet is known as a private subnet.

If a subnet doesn't have a route to the internet gateway, but has its traffic routed to a virtual private gateway, the subnet is known as a VPN-only subnet.

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

#### NEW QUESTION 24

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 28

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

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

**Answer: B**

#### Explanation:

Auto Scaling enables you to follow the demand curve for your applications closely, reducing the need to manually provision Amazon EC2 capacity in advance. For example, you can set a condition to add new

Amazon EC2 instances in increments to the Auto Scaling group when the average utilization of your Amazon EC2 fleet is high; and similarly, you can set a condition to remove instances in the same increments when CPU utilization is low. If you have predictable load changes, you can set a schedule through Auto Scaling to plan your scaling activities. You can use Amazon CloudWatch to send alarms to trigger scaling activities and Elastic Load Balancing to help distribute traffic to your instances within Auto Scaling groups. Auto Scaling enables you to run your Amazon EC2 fleet at optimal utilization. Reference:

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

#### NEW QUESTION 31

You are building infrastructure for a data warehousing solution and an extra request has come through that there will be a lot of business reporting queries running all the time and you are not sure if your current DB instance will be able to handle it. What would be the best solution for this?

- A. DB Parameter Groups
- B. Read Replicas
- C. Multi-AZ DB Instance deployment
- D. Database Snapshots

**Answer: B**

#### Explanation:

Read Replicas make it easy to take advantage of MySQL's built-in replication functionality to elastically scale out beyond the capacity constraints of a single DB Instance for read-heavy database workloads. There are a variety of scenarios where deploying one or more Read Replicas for a given source DB Instance may make sense. Common reasons for deploying a Read Replica include:

Scaling beyond the compute or I/O capacity of a single DB Instance for read-heavy database workloads. This excess read traffic can be directed to one or more Read Replicas.

Serving read traffic while the source DB Instance is unavailable. If your source DB Instance cannot take I/O requests (e.g. due to I/O suspension for backups or scheduled maintenance), you can direct read traffic to your Read Replica(s). For this use case, keep in mind that the data on the Read Replica may be "stale" since the source DB Instance is unavailable.

Business reporting or data warehousing scenarios; you may want business reporting queries to run against a Read Replica, rather than your primary, production DB Instance.

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

#### NEW QUESTION 32

Much of your company's data does not need to be accessed often, and can take several hours for retrieval time, so it's stored on Amazon Glacier. However someone within your organization has expressed concerns that his data is more sensitive than the other data, and is wondering whether the high

level of encryption that he knows is on S3 is also used on the much cheaper Glacier service. Which of the following statements would be most applicable in regards to this concern?

- A. There is no encryption on Amazon Glacier, that's why it is cheaper.
- B. Amazon Glacier automatically encrypts the data using AES-128 a lesser encryption method than Amazon S3 but you can change it to AES-256 if you are willing to pay more.
- C. Amazon Glacier automatically encrypts the data using AES-256, the same as Amazon S3.
- D. Amazon Glacier automatically encrypts the data using AES-128 a lesser encryption method than Amazon S3.

**Answer: C**

**Explanation:**

Like Amazon S3, the Amazon Glacier service provides low-cost, secure, and durable storage. But where S3 is designed for rapid retrieval, Glacier is meant to be used as an archival service for data that is not accessed often, and for which retrieval times of several hours are suitable.

Amazon Glacier automatically encrypts the data using AES-256 and stores it durably in an immutable form. Amazon Glacier is designed to provide average annual durability of 99.999999999% for an archive. It stores each archive in multiple facilities and multiple devices. Unlike traditional systems which can require laborious data verification and manual repair, Glacier performs regular, systematic data integrity checks, and is built to be automatically self-healing.

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

**NEW QUESTION 36**

In Amazon RDS, security groups are ideally used to:

- A. Define maintenance period for database engines
- B. Launch Amazon RDS instances in a subnet
- C. Create, describe, modify, and delete DB instances
- D. Control what IP addresses or EC2 instances can connect to your databases on a DB instance

**Answer: D**

**Explanation:**

In Amazon RDS, security groups are used to control what IP addresses or EC2 instances can connect to your databases on a DB instance.

When you first create a DB instance, its firewall prevents any database access except through rules specified by an associated security group.

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

**NEW QUESTION 41**

You need to set up a complex network infrastructure for your organization that will be reasonably easy to deploy, replicate, control, and track changes on. Which AWS service would be best to use to help you accomplish this?

- A. AWS Import/Export
- B. AWS CloudFormation
- C. Amazon Route 53
- D. Amazon CloudWatch

**Answer: B**

**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. You don't need to manually create and configure AWS resources and figure out what's dependent on what. AWS CloudFormation handles all of that.

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

**NEW QUESTION 44**

You have just been given a scope for a new client who has an enormous amount of data (petabytes) that he constantly needs analysed. Currently he is paying a huge amount of money for a data warehousing company to do this for him and is wondering if AWS can provide a cheaper solution. Do you think AWS has a solution for this?

- A. Yes
- B. Amazon SimpleDB
- C. No
- D. Not presently
- E. Yes
- F. Amazon Redshift
- G. Yes
- H. Your choice of relational AMIs on Amazon EC2 and EBS

**Answer: C**

**Explanation:**

Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse service that makes it simple and cost-effective to efficiently analyze all your data using your existing business intelligence tools. You can start small for just \$0.25 per hour with no commitments or upfront costs and scale to a petabyte or more for \$1,000 per terabyte per year, less than a tenth of most other data warehousing solutions. Amazon Redshift delivers fast query performance by using columnar storage technology to improve I/O efficiency and parallelizing queries across multiple nodes. Redshift uses standard PostgreSQL JDBC and ODBC drivers, allowing you to use a wide range of familiar SQL clients. Data load speed scales linearly with cluster size, with integrations to Amazon S3, Amazon DynamoDB, Amazon Elastic MapReduce, Amazon Kinesis or any SSH-enabled host.

Reference: [https://aws.amazon.com/running\\_databases/#redshift\\_anchor](https://aws.amazon.com/running_databases/#redshift_anchor)

**NEW QUESTION 45**

In Amazon EC2, while sharing an Amazon EBS snapshot, can the snapshots with AWS Marketplace product codes be public?

- A. Yes, but only for US-based providers.
- B. Yes, they can be public.
- C. No, they cannot be made public.
- D. Yes, they are automatically made public by the system

**Answer: C**

**Explanation:**

Snapshots with AWS Marketplace product codes can't be made public. Reference:  
<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/ebs-modifying-snapshot-permissions.html>

**NEW QUESTION 50**

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 53**

You are running PostgreSQL on Amazon RDS and it seems to be all running smoothly deployed in one availability zone. A database administrator asks you if DB instances running PostgreSQL support Multi-AZ deployments. What would be a correct response to this QUESTION ?

- A. Yes.
- B. Yes but only for small db instances.
- C. No.
- D. Yes but you need to request the service from AWS

**Answer: A**

**Explanation:**

Amazon RDS supports DB instances running several versions of PostgreSQL. Currently we support PostgreSQL versions 9.3.1, 9.3.2, and 9.3.3. You can create DB instances and DB snapshots, point-in-time restores and backups.

DB instances running PostgreSQL support Multi-AZ deployments, Provisioned IOPS, and can be created inside a VPC. You can also use SSL to connect to a DB instance running PostgreSQL.

You can use any standard SQL client application to run commands for the instance from your client computer. Such applications include pgAdmin, a popular Open Source administration and development tool for PostgreSQL, or psql, a command line utility that is part of a PostgreSQL installation. In order to deliver a managed service experience, Amazon RDS does not provide host access to DB instances, and it restricts access to certain system procedures and tables that require advanced privileges. Amazon RDS supports access to databases on a DB instance using any standard SQL client application. Amazon RDS does not allow direct host access to a DB instance via Telnet or Secure Shell (SSH).

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

**NEW QUESTION 57**

A user has launched 10 EC2 instances inside a placement group. Which of the below mentioned statements is true with respect to the placement group?

- A. All instances must be in the same AZ
- B. All instances can be across multiple regions
- C. The placement group cannot have more than 5 instances
- D. All instances must be in the same region

**Answer: A**

**Explanation:**

A placement group is a logical grouping of EC2 instances within a single Availability Zone. Using placement groups enables applications to participate in a low-latency, 10 Gbps network. Placement groups are recommended for applications that benefit from low network latency, high network throughput or both.

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

**NEW QUESTION 58**

An organization has developed a mobile application which allows end users to capture a photo on their mobile device, and store it inside an application. The application internally uploads the data to AWS S3. The organization wants each user to be able to directly upload data to S3 using their Google ID. How will the mobile app allow this?

- A. Use the AWS Web identity federation for mobile applications, and use it to generate temporary security credentials for each user.
- B. It is not possible to connect to AWS S3 with a Google ID.
- C. Create an IAM user every time a user registers with their Google ID and use IAM to upload files to S3.
- D. Create a bucket policy with a condition which allows everyone to upload if the login ID has a Google part to it.

**Answer: A**

**Explanation:**

For Amazon Web Services, the Web identity federation allows you to create cloud-backed mobile apps that use public identity providers, such as login with Facebook, Google, or Amazon. It will create temporary security credentials for each user, which will be authenticated by the AWS services, such as S3.  
Reference: <http://docs.aws.amazon.com/STS/latest/UsingSTS/CreatingWIF.html>

**NEW QUESTION 62**

You are architecting an auto-scalable batch processing system using video processing pipelines and Amazon Simple Queue Service (Amazon SQS) for a customer. You are unsure of the limitations of SQS and need to find out. What do you think is a correct statement about the limitations of Amazon SQS?

- A. It supports an unlimited number of queues but a limited number of messages per queue for each user but automatically deletes messages that have been in the queue for more than 4 weeks.
- B. It supports an unlimited number of queues and unlimited number of messages per queue for each user but automatically deletes messages that have been in the queue for more than 4 days.
- C. It supports an unlimited number of queues but a limited number of messages per queue for each user but automatically deletes messages that have been in the queue for more than 4 days.
- D. It supports an unlimited number of queues and unlimited number of messages per queue for each user but automatically deletes messages that have been in the queue for more than 4 weeks.

**Answer: B**

**Explanation:**

Amazon Simple Queue Service (Amazon SQS) is a messaging queue service that handles message or workflows between other components in a system. Amazon SQS supports an unlimited number of queues and unlimited number of messages per queue for each user. Please be aware that Amazon SQS automatically deletes messages that have been in the queue for more than 4 days.  
Reference: <http://aws.amazon.com/documentation/sqs/>

**NEW QUESTION 63**

Once again your customers are concerned about the security of their sensitive data and with their latest enquiry ask about what happens to old storage devices on AWS. What would be the best answer to this QUESTION ?

- A. AWS reformats the disks and uses them again.
- B. AWS uses the techniques detailed in DoD 5220.22-M to destroy data as part of the decommissioning process.
- C. AWS uses their own proprietary software to destroy data as part of the decommissioning process.
- D. AWS uses a 3rd party security organization to destroy data as part of the decommissioning process

**Answer: B**

**Explanation:**

When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals.  
AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual ") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process.  
All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.  
Reference: <http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf>

**NEW QUESTION 65**

Your company has been storing a lot of data in Amazon Glacier and has asked for an inventory of what is in there exactly. So you have decided that you need to download a vault inventory. Which of the following statements is incorrect in relation to Vault Operations in Amazon Glacier?

- A. You can use Amazon Simple Notification Service (Amazon SNS) notifications to notify you when the job completes.
- B. A vault inventory refers to the list of archives in a vault.
- C. You can use Amazon Simple Queue Service (Amazon SQS) notifications to notify you when the job completes.
- D. Downloading a vault inventory is an asynchronous operation

**Answer: C**

**Explanation:**

Amazon Glacier supports various vault operations.  
A vault inventory refers to the list of archives in a vault. For each archive in the list, the inventory provides archive information such as archive ID, creation date, and size. Amazon Glacier updates the vault inventory approximately once a day, starting on the day the first archive is uploaded to the vault. A vault inventory must exist for you to be able to download it.  
Downloading a vault inventory is an asynchronous operation. You must first initiate a job to download the inventory. After receiving the job request, Amazon Glacier prepares your inventory for download. After the job completes, you can download the inventory data.  
Given the asynchronous nature of the job, you can use Amazon Simple Notification Service (Amazon SNS) notifications to notify you when the job completes. You can specify an Amazon SNS topic for each individual job request or configure your vault to send a notification when specific vault events occur. Amazon Glacier prepares an inventory for each vault periodically, every 24 hours. If there have been no archive additions or deletions to the vault since the last inventory, the inventory date is not updated. When you initiate a job for a vault inventory, Amazon Glacier returns the last inventory it generated, which is a point-in-time snapshot and not real-time data. You might not find it useful to retrieve vault inventory for each archive upload. However, suppose you maintain a database on the client-side associating metadata about the archives you upload to Amazon Glacier. Then, you might find the vault inventory useful to reconcile information in your database with the actual vault inventory.  
Reference: <http://docs.aws.amazon.com/amazonglacier/latest/dev/working-with-vaults.html>

**NEW QUESTION 67**

A customer enquires about whether all his data is secure on AWS and is especially concerned about Elastic Map Reduce (EMR) so you need to inform him of some of the security features in place for AWS. Which of the below statements would be an incorrect response to your customers enquiry?

- A. Amazon EMR customers can choose to send data to Amazon S3 using the HTTPS protocol for secure transmission.
- B. Amazon S3 provides authentication mechanisms to ensure that stored data is secured against unauthorized access.

- C. Every packet sent in the AWS network uses Internet Protocol Security (IPsec).  
D. Customers may encrypt the input data before they upload it to Amazon S3.

**Answer: C**

**Explanation:**

Amazon S3 provides authentication mechanisms to ensure that stored data is secured against unauthorized access. Unless the customer who is uploading the data specifies otherwise, only that customer can access the data. Amazon EMR customers can also choose to send data to Amazon S3 using the HTTPS protocol for secure transmission. In addition, Amazon EMR always uses HTTPS to send data between Amazon S3 and Amazon EC2. For added security, customers may encrypt the input data before they upload it to Amazon S3 (using any common data compression tool); they then need to add a decryption step to the beginning of their cluster when Amazon EMR fetches the data from Amazon S3. Reference: <https://aws.amazon.com/elasticmapreduce/faqs/>

**NEW QUESTION 69**

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 instance.

**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 74**

What would be the best way to retrieve the public IP address of your EC2 instance using the CLI?

- A. Using tags  
B. Using traceroute  
C. Using ipconfig  
D. Using instance metadata

**Answer: D**

**Explanation:**

To determine your instance's public IP address from within the instance, you can use instance metadata. Use the following command to access the public IP address: For Linux use, `$ curl`

`http://169.254.169.254/latest/meta-data/public-ipv4`, and for Windows use, `$ wget http://169.254.169.254/latest/meta-data/public-ipv4`.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-instance-addressing.htm>

**NEW QUESTION 75**

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 77**

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 81**

A user has configured ELB with two EBS backed EC2 instances. The user is trying to understand the DNS access and IP support for ELB. Which of the below mentioned statements may not help the user understand the IP mechanism supported by ELB?

- A. The client can connect over IPV4 or IPV6 using Dualstack
- B. Communication between the load balancer and back-end instances is always through IPV4
- C. ELB DNS supports both IPV4 and IPV6
- D. The ELB supports either IPV4 or IPV6 but not both

**Answer: D**

**Explanation:**

Elastic Load Balancing supports both Internet Protocol version 6 (IPv6) and Internet Protocol version 4 (IPv4). Clients can connect to the user's load balancer using either IPv4 or IPv6 (in EC2-Classic) DNS. However, communication between the load balancer and its back-end instances uses only IPv4. The user can use the Dualstack-prefixed DNS name to enable IPv6 support for communications between the client and the load balancers. Thus, the clients are able to access the load balancer using either IPv4 or IPv6 as their individual connectivity needs dictate.

Reference: <http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/UserScenariosForEC2.html>

**NEW QUESTION 85**

Amazon S3 allows you to set per-file permissions to grant read and/or write access. However you have decided that you want an entire bucket with 100 files already in it to be accessible to the public. You don't want to go through 100 files individually and set permissions. What would be the best way to do this?

- A. Move the bucket to a new region
- B. Add a bucket policy to the bucket.
- C. Move the files to a new bucket.
- D. Use Amazon EBS instead of S3

**Answer: B**

**Explanation:**

Amazon S3 supports several mechanisms that give you flexibility to control who can access your data as well as how, when, and where they can access it. Amazon S3 provides four different access control mechanisms: AWS Identity and Access Management (IAM) policies, Access Control Lists (ACLs), bucket policies, and query string authentication. IAM enables organizations to create and manage multiple users under a single AWS account. With IAM policies, you can grant IAM users fine-grained control to your Amazon S3 bucket or objects. You can use ACLs to selectively add (grant) certain permissions on individual objects. Amazon S3 bucket policies can be used to add or deny permissions across some or all of the objects within a single bucket. With Query string authentication, you have the ability to share Amazon S3 objects through URLs that are valid for a specified period of time.

Reference: <http://aws.amazon.com/s3/details/#security>

**NEW QUESTION 89**

You need to set up a high level of security for an Amazon Relational Database Service (RDS) you have just built in order to protect the confidential information stored in it. What are all the possible security groups that RDS uses?

- A. DB security groups, VPC security groups, and EC2 security groups.
- B. DB security groups only.
- C. EC2 security groups only.
- D. VPC security groups, and EC2 security group

**Answer: A**

**Explanation:**

A security group controls the access to a DB instance. It does so by allowing access to IP address ranges or Amazon EC2 instances that you specify. Amazon RDS uses DB security groups, VPC security groups, and EC2 security groups. In simple terms, a DB security group controls access to a DB instance that is not in a VPC, a VPC security group controls access to a DB instance inside a VPC, and an Amazon EC2 security group controls access to an EC2 instance and can be used with a DB instance.

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

**NEW QUESTION 92**

You have set up an Elastic Load Balancer (ELB) with the usual default settings, which route each request independently to the application instance with the smallest load. However, someone has asked you to bind a user's session to a specific application instance so as to ensure that all requests coming from the user during the session will be sent to the same application instance. AWS has a feature to do this. What is it called?

- A. Connection draining
- B. Proxy protocol
- C. Tagging
- D. Sticky session

**Answer: D**

**Explanation:**

An Elastic Load Balancer(ELB) by default, routes each request independently to the application instance with the smallest load. However, you can use the sticky session feature (also known as session affinity), which enables the load balancer to bind a user's session to a specific application instance. This ensures that all requests coming from the user during the session will be sent to the same application instance. The key to managing the sticky session is determining how long your load balancer should consistently route the user's request to the same application instance. If your application has its own session cookie, then you can set Elastic Load Balancing to create the session cookie to follow the duration specified by the application's

session cookie. If your application does not have its own session cookie, then you can set Elastic Load Balancing to create a session cookie by specifying your own stickiness duration. You can associate stickiness duration for only HTTP/HTTPS load balancer listeners.

An application instance must always receive and send two cookies: A cookie that defines the stickiness duration and a special Elastic Load Balancing cookie named AWSELB, that has the mapping to the application instance.

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

**NEW QUESTION 94**

A user has created an application which will be hosted on EC2. The application makes calls to DynamoDB to fetch certain data. The application is using the DynamoDB SDK to connect with from the EC2 instance. Which of the below mentioned statements is true with respect to the best practice for security in this scenario?

- A. The user should create an IAM user with DynamoDB access and use its credentials within the application to connect with DynamoDB
- B. The user should attach an IAM role with DynamoDB access to the EC2 instance
- C. The user should create an IAM role, which has EC2 access so that it will allow deploying the application
- D. The user should create an IAM user with DynamoDB and EC2 access
- E. Attach the user with the application so that it does not use the root account credentials

**Answer: B**

**Explanation:**

With AWS IAM a user is creating an application which runs on an EC2 instance and makes requests to AWS, such as DynamoDB or S3 calls. Here it is recommended that the user should not create an IAM user and pass the user's credentials to the application or embed those credentials inside the application. Instead, the user should use roles for EC2 and give that role access to DynamoDB /S3. When the roles are attached to EC2, it will give temporary security credentials to the application hosted on that EC2, to connect with DynamoDB / S3.

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

**NEW QUESTION 99**

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 100**

A user is aware that a huge download is occurring on his instance. He has already set the Auto Scaling policy to increase the instance count when the network I/O increases beyond a certain limit. How can the user ensure that this temporary event does not result in scaling?

- A. The network I/O are not affected during data download
- B. The policy cannot be set on the network I/O
- C. There is no way the user can stop scaling as it is already configured
- D. Suspend scaling

**Answer: D**

**Explanation:**

The user may want to stop the automated scaling processes on the Auto Scaling groups either to perform manual operations or during emergency situations. To perform this, the user can suspend one or more scaling processes at any time. Once it is completed, the user can resume all the suspended processes.

Reference: [http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS\\_Concepts.html](http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html)

**NEW QUESTION 103**

Which of the following strategies can be used to control access to your Amazon EC2 instances?

- A. DB security groups
- B. IAM policies
- C. None of these
- D. EC2 security groups

**Answer: D**

**Explanation:**

IAM policies allow you to specify what actions your IAM users are allowed to perform against your EC2 Instances. However, when it comes to access control, security groups are what you need in order to define and control the way you want your instances to be accessed, and whether or not certain kind of communications are allowed or not.

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

**NEW QUESTION 106**

A client of yours has a huge amount of data stored on Amazon S3, but is concerned about someone stealing it while it is in transit. You know that all data is encrypted in transit on AWS, but which of the following is wrong when describing server-side encryption on AWS?

- A. Amazon S3 server-side encryption employs strong multi-factor encryption.
- B. Amazon S3 server-side encryption uses one of the strongest block ciphers available, 256-bit Advanced Encryption Standard (AES-256), to encrypt your data.
- C. In server-side encryption, you manage encryption/decryption of your data, the encryption keys, and related tools.
- D. Server-side encryption is about data encryption at rest—that is, Amazon S3 encrypts your data as it writes it to disks.

**Answer: C**

**Explanation:**

Amazon S3 encrypts your object before saving it on disks in its data centers and decrypts it when you download the objects. You have two options depending on how you choose to manage the encryption keys: Server-side encryption and client-side encryption.

Server-side encryption is about data encryption at rest—that is, Amazon S3 encrypts your data as it writes it to disks in its data centers and decrypts it for you when you access it. As long as you authenticate your request and you have access permissions, there is no difference in the way you access encrypted or unencrypted objects. Amazon S3 manages encryption and decryption for you. For example, if you share your objects using a pre-signed URL, that URL works the same way for both encrypted and unencrypted objects.

In client-side encryption, you manage encryption/decryption of your data, the encryption keys, and related tools. Server-side encryption is an alternative to client-side encryption in which Amazon S3 manages the encryption of your data, freeing you from the tasks of managing encryption and encryption keys.

Amazon S3 server-side encryption employs strong multi-factor encryption. Amazon S3 encrypts each object with a unique key. As an additional safeguard, it encrypts the key itself with a master key that it regularly rotates. Amazon S3 server-side encryption uses one of the strongest block ciphers available, 256-bit Advanced Encryption Standard (AES-256), to encrypt your data.

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

**NEW QUESTION 108**

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 113**

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 114**

You have a number of image files to encode. In an Amazon SQS worker queue, you create an Amazon SQS message for each file specifying the command (jpeg-encode) and the location of the file in Amazon S3. Which of the following statements best describes the functionality of Amazon SQS?

- A. Amazon SQS is a distributed queuing system that is optimized for horizontal scalability, not for single-threaded sending or receive speeds.
- B. Amazon SQS is for single-threaded sending or receive speeds.
- C. Amazon SQS is a non-distributed queuing system.
- D. Amazon SQS is a distributed queuing system that is optimized for vertical scalability and for single-threaded sending or receive speeds.

**Answer: A**

**Explanation:**

Amazon SQS is a distributed queuing system that is optimized for horizontal scalability, not for single-threaded sending or receive speeds. A single client can send or receive Amazon SQS messages at a rate of about 5 to 50 messages per second. Higher receive performance can be achieved by requesting multiple messages (up to 10) in a single call. It may take several seconds before a message that has been to a queue is available to be received.

Reference: [http://media.amazonwebservices.com/AWS\\_Storage\\_Options.pdf](http://media.amazonwebservices.com/AWS_Storage_Options.pdf)

**NEW QUESTION 117**

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 118**

A user has set up the CloudWatch alarm on the CPU utilization metric at 50%, with a time interval of 5 minutes and 10 periods to monitor. What will be the state of the alarm at the end of 90 minutes, if the CPU utilization is constant at 80%?

- A. ALERT
- B. ALARM
- C. OK
- D. INSUFFICIENT\_DATA

**Answer: B**

**Explanation:**

In this case the alarm watches a metric every 5 minutes for 10 intervals. Thus, it needs at least 50 minutes to come to the "OK" state. Till then it will be in the INSUFFICIENT\_DATA state.

Since 90 minutes have passed and CPU utilization is at 80% constant, the state of alarm will be "ALARM". Reference: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/AlarmThatSendsEmail.html>

**NEW QUESTION 123**

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

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

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

**Answer: B**

**Explanation:**

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

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

**NEW QUESTION 127**

Which one of the below is not an AWS Storage Service?

- A. Amazon S3
- B. Amazon Glacier
- C. Amazon CloudFront
- D. Amazon EBS

**Answer: C**

**Explanation:**

AWS Storage Services are: Amazon S3

Amazon Glacier Amazon EBS

AWS Storage Gateway

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

**NEW QUESTION 129**

Regarding Amazon Route 53, 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 to route traffic to the correct region and then use to route traffic to instances within the region, based on probabilities that you specify.

- A. weighted-based routing; alias resource record sets
- B. latency-based routing; weighted resource record sets
- C. weighted-based routing; weighted resource record sets
- D. latency-based routing; alias resource record sets

**Answer: B**

**Explanation:**

Regarding Amazon Route 53, 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.

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

**NEW QUESTION 130**

You receive a bill from AWS but are confused because you see you are incurring different costs for the exact same storage size in different regions on Amazon S3. You ask AWS why this is so. What response would you expect to receive from AWS?

- A. We charge less in different time zones.
- B. We charge less where our costs are less.
- C. This will balance out next bill.
- D. It must be a mistake

**Answer: B**

**Explanation:**

Amazon S3 is storage for the internet. It's a simple storage service that offers software developers a highly-scalable, reliable, and low-latency data storage infrastructure at very low costs.

AWS charges less where their costs are less.

For example, their costs are lower in the US Standard Region than in the US West (Northern California) Region.

Reference: <https://aws.amazon.com/s3/faqs/>

**NEW QUESTION 134**

You are setting up some EBS volumes for a customer who has requested a setup which includes a RAID (redundant array of inexpensive disks). AWS has some recommendations for RAID setups. Which RAID setup is not recommended for Amazon EBS?

- A. RAID 5 only
- B. RAID 5 and RAID 6
- C. RAID 1 only
- D. RAID 1 and RAID 6

**Answer: B**

**Explanation:**

With Amazon EBS, you can use any of the standard RAID configurations that you can use with a traditional bare metal server, as long as that particular RAID configuration is supported by the operating system for your instance. This is because all RAID is accomplished at the software level. For greater I/O performance than you can achieve with a single volume, RAID 0 can stripe multiple volumes together; for on-instance redundancy, RAID 1 can mirror two volumes together. RAID 5 and RAID 6 are not recommended for Amazon EBS because the parity write operations of these RAID modes consume some of the IOPS available to your volumes.

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

**NEW QUESTION 139**

A user is hosting a website in the US West-1 region. The website has the highest client base from the Asia-Pacific (Singapore / Japan) region. The application is accessing data from S3 before serving it to client. Which of the below mentioned regions gives a better performance for S3 objects?

- A. Japan
- B. Singapore
- C. US East
- D. US West-1

**Answer: D**

**Explanation:**

Access to Amazon S3 from within Amazon EC2 in the same region is fast. In this aspect, though the client base is Singapore, the application is being hosted in the US West-1 region. Thus, it is recommended that S3 objects be stored in the US-West-1 region.

Reference: [http://media.amazonwebservices.com/AWS\\_Storage\\_Options.pdf](http://media.amazonwebservices.com/AWS_Storage_Options.pdf)

**NEW QUESTION 144**

The AWS CloudHSM service defines a resource known as a high-availability (HA) , which is a virtual partition that represents a group of partitions, typically distributed between several physical HSMs for high-availability.

- A. proxy group
- B. partition group
- C. functional group
- D. relational group

**Answer: B**

**Explanation:**

The AWS CloudHSM service defines a resource known as a high-availability (HA) partition group, which is a virtual partition that represents a group of partitions, typically distributed between several physical HSMs for high-availability.

Reference: <http://docs.aws.amazon.com/cloudhsm/latest/userguide/configuring-ha.html>

**NEW QUESTION 146**

Is it possible to get a history of all EC2 API calls made on your account for security analysis and operational troubleshooting purposes?

- A. Yes, by default, the history of your API calls is logged.
- B. Yes, you should turn on the CloudTrail in the AWS console.
- C. No, you can only get a history of VPC API calls.
- D. No, you cannot store history of EC2 API calls on Amazon.

**Answer: B**

**Explanation:**

To get a history of all EC2 API calls (including VPC and EBS) made on your account, you simply turn on CloudTrail in the AWS Management Console.

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

#### NEW QUESTION 148

You have just set up your first Elastic Load Balancer (ELB) but it does not seem to be configured properly. You discover that before you start using ELB, you have to configure the listeners for your load balancer. Which protocols does ELB use to support the load balancing of applications?

- A. HTTP and HTTPS
- B. HTTP, HTTPS, TCP, SSL and SSH
- C. HTTP, HTTPS, TCP, and SSL
- D. HTTP, HTTPS, TCP, SSL and SFTP

**Answer:** C

#### Explanation:

Before you start using Elastic Load Balancing (ELB), you have to configure the listeners for your load balancer. A listener is a process that listens for connection requests. It is configured with a protocol and a port number for front-end (client to load balancer) and back-end (load balancer to back-end instance) connections. Elastic Load Balancing supports the load balancing of applications using HTTP, HTTPS (secure HTTP), TCP, and SSL (secure TCP) protocols. The HTTPS uses the SSL protocol to establish secure connections over the HTTP layer. You can also use SSL protocol to establish secure connections over the TCP layer. The acceptable ports for both HTTPS/SSL and HTTP/TCP connections are 25, 80, 443, 465, 587, and 1024-65535.

Reference:

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

#### NEW QUESTION 151

When does the billing of an Amazon EC2 system begin?

- A. It starts when the Status column for your distribution changes from Creating to Deployed.
- B. It starts as soon as you click the create instance option on the main EC2 console.
- C. It starts when your instance reaches 720 instance hours.
- D. It starts when Amazon EC2 initiates the boot sequence of an AM instance.

**Answer:** D

#### Explanation:

Billing commences when Amazon EC2 initiates the boot sequence of an AM instance. Billing ends when the instance terminates, which could occur through a web services command, by running "shutdown -h", or through instance failure. When you stop an instance, Amazon shuts it down but doesn't charge hourly usage for a stopped instance, or data transfer fees, but charges for the storage for any Amazon EBS volumes.

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

#### NEW QUESTION 152

What is the data model of DynamoDB?

- A. Since DynamoDB is schema-less, there is no data model.
- B. "Items", with Keys and one or more Attribute; and "Attribute", with Name and Value.
- C. "Table", a collection of Items; "Items", with Keys and one or more Attribute; and "Attribute", with Name and Value.
- D. "Database", which is a set of "Tables", which is a set of "Items", which is a set of "Attributes".

**Answer:** C

#### Explanation:

The data model of DynamoDB is: "Table", a collection of Items;

"Items", with Keys and one or more Attribute; "Attribute", with Name and Value.

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

#### NEW QUESTION 154

What happens to Amazon EBS root device volumes, by default, when an instance terminates?

- A. Amazon EBS root device volumes are moved to IAM.
- B. Amazon EBS root device volumes are copied into Amazon RDS.
- C. Amazon EBS root device volumes are automatically deleted.
- D. Amazon EBS root device volumes remain in the database until you delete the

**Answer:** C

#### Explanation:

By default, Amazon EBS root device volumes are automatically deleted when the instance terminates. Reference:

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

#### NEW QUESTION 159

A gaming company comes to you and asks you to build their infrastructure for their site. They are not sure how big they will be as with all start-ups they have limited money and big ideas. What they do tell you is that if the game becomes successful, like one of their previous games, it may rapidly grow to millions of users and generate tens (or even hundreds) of thousands of writes and reads per second. After considering all of this, you decide that they need a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. Which of the following databases do you think would best fit their needs?

- A. Amazon DynamoDB
- B. Amazon Redshift
- C. Any non-relational database.
- D. Amazon SimpleDB

**Answer:** A

**Explanation:**

Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. Amazon DynamoDB enables customers to offload the administrative burdens of operating and scaling distributed databases to AWS, so they don't have to worry about hardware provisioning, setup and configuration, replication, software patching, or cluster scaling. Today's web-based applications generate and consume massive amounts of data. For example, an online game might start out with only a few thousand users and a light database workload consisting of 10 writes per second and 50 reads per second. However, if the game becomes successful, it may rapidly grow to millions of users and generate tens (or even hundreds) of thousands of writes and reads per second. It may also create terabytes or more of data per day. Developing your applications against Amazon DynamoDB enables you to start small and simply dial-up your request capacity for a table as your requirements scale, without incurring downtime. You pay highly cost-efficient rates for the request capacity you provision, and let Amazon DynamoDB do the work over partitioning your data and traffic over sufficient server capacity to meet your needs. Amazon DynamoDB does the database management and administration, and you simply store and request your data. Automatic replication and failover provides built-in fault tolerance, high availability, and data durability. Amazon DynamoDB gives you the peace of mind that your database is fully managed and can grow with your application requirements. Reference: <http://aws.amazon.com/dynamodb/faqs/>

**NEW QUESTION 160**

A favored client needs you to quickly deploy a database that is a relational database service with minimal administration as he wants to spend the least amount of time administering it. Which database would be the best option?

- A. Amazon SimpleDB
- B. Your choice of relational AMs on Amazon EC2 and EBS.
- C. Amazon RDS
- D. Amazon Redshift

**Answer:** C

**Explanation:**

Amazon Relational Database Service (Amazon RDS) is a web 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.

Amazon RDS gives you access to the capabilities of a familiar MySQL, Oracle, SQL Server, or PostgreSQL database engine. This means that the code, applications, and tools you already use today with your existing databases can be used with Amazon RDS. Amazon RDS automatically patches the database software and backs up your database, storing the backups for a user-defined retention period and enabling point-in-time recovery.

Reference: [https://aws.amazon.com/running\\_databases/#rds\\_anchor](https://aws.amazon.com/running_databases/#rds_anchor)

**NEW QUESTION 162**

You're trying to delete an SSL certificate from the IAM certificate store, and you're getting the message "Certificate: <certificate-id> is being used by CloudFront." Which of the following statements is probably the reason why you are getting this error?

- A. Before you can delete an SSL certificate, you need to either rotate SSL certificates or revert from using a custom SSL certificate to using the default CloudFront certificate.
- B. You can't delete SSL certificates. You need to request it from AWS.
- C. Before you can delete an SSL certificate, you need to set up the appropriate access level in IAM
- D. Before you can delete an SSL certificate you need to set up https on your server

**Answer:** A

**Explanation:**

CloudFront is a web service that speeds up distribution of your static and dynamic web content, for example, .html, .css, .php, and image files, to end users. Every CloudFront web distribution must be associated either with the default CloudFront certificate or with a custom SSL certificate. Before you can delete an SSL certificate, you need to either rotate SSL certificates (replace the current custom SSL certificate with another custom SSL certificate) or revert from using a custom SSL certificate to using the default CloudFront certificate.

Reference: <http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Troubleshooting.html>

**NEW QUESTION 165**

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 168**

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 170**

In relation to AWS CloudHSM, High-availability (HA) recovery is hands-off resumption by failed HA group members.

Prior to the introduction of this function, the HA feature provided redundancy and performance, but required that a failed/lost group member be reinstated.

- A. automatically
- B. periodically
- C. manually
- D. continuously

**Answer: C**

**Explanation:**

In relation to AWS CloudHSM, High-availability (HA) recovery is hands-off resumption by failed HA group members.

Prior to the introduction of this function, the HA feature provided redundancy and performance, but required that a failed/lost group member be manually reinstated.

Reference: <http://docs.aws.amazon.com/cloudhsm/latest/userguide/ha-best-practices.html>

**NEW QUESTION 171**

You have created a Route 53 latency record set from your domain to a machine in Northern Virginia and a similar record to a machine in Sydney.

When a user located in U S visits your domain he will be routed to:

- A. Northern Virginia
- B. Sydney
- C. Both, Northern Virginia and Sydney
- D. Depends on the Weighted Resource Record Sets

**Answer: A**

**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 173**

Having just set up your first Amazon Virtual Private Cloud (Amazon VPC) network, which defined a default network interface, you decide that you need to create and attach an additional network interface, known as an elastic network interface (ENI) to one of your instances. Which of the following statements is true regarding attaching network interfaces to your instances in your VPC?

- A. You can attach 5 ENIs per instance type.
- B. You can attach as many ENIs as you want.
- C. The number of ENIs you can attach varies by instance type.
- D. You can attach 100 ENIs total regardless of instance type.

**Answer: C**

**Explanation:**

Each instance in your VPC has a default network interface that is assigned a private IP address from the IP address range of your VPC. You can create and attach an additional network interface, known as an elastic network interface (ENI), to any instance in your VPC. The number of ENIs you can attach varies by instance type.

**NEW QUESTION 178**

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 Multi-AZ one.

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

**NEW QUESTION 179**

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 183**

You have been asked to build AWS infrastructure for disaster recovery for your local applications and within that you should use an AWS Storage Gateway as part of the solution. Which of the following best describes the function of an AWS Storage Gateway?

- A. Accelerates transferring large amounts of data between the AWS cloud and portable storage devices .
- B. A web service that speeds up distribution of your static and dynamic web content.
- C. Connects an on-premises software appliance with cloud-based storage to provide seamless and secure integration between your on-premises IT environment and AWS's storage infrastructure.
- D. Is a storage service optimized for infrequently used data, or "cold data."

**Answer: C**

**Explanation:**

AWS Storage Gateway connects an on-premises software appliance with cloud-based storage to provide seamless integration with data security features between your on-premises IT environment and the Amazon Web Services (AWS) storage infrastructure. You can use the service to store data in the AWS cloud for scalable and cost-effective storage that helps maintain data security. AWS Storage Gateway offers both volume-based and tape-based storage solutions:

Volume gateways Gateway-cached volumes Gateway-stored volumes

Gateway-virtual tape library (VTL)

Reference:

[http://media.amazonwebservices.com/architecturecenter/AWS\\_ac\\_ra\\_disasterrecovery\\_07.pdf](http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_disasterrecovery_07.pdf)

**NEW QUESTION 187**

Which of the following features are provided by Amazon EC2?

- A. Exadata Database Machine, Optimized Storage Management, Flashback Technology, and Data Warehousing
- B. Instances, Amazon Machine Images (AMIs), Key Pairs, Amazon EBS Volumes, Firewall, Elastic IP address, Tags, and Virtual Private Clouds (VPCs)
- C. Real Application Clusters (RAC), ElastiCache Machine Images (EMIs), Data Warehousing, Flashback Technology, Dynamic IP address
- D. Exadata Database Machine, Real Application Clusters (RAC), Data Guard, Table and Index Partitioning, and Data Pump Compression

**Answer: B**

**Explanation:**

Amazon EC2 provides the following features:

- Virtual computing environments, known as instances;
- Pre-configured templates for your instances, known as Amazon Machine Images (AMIs), that package the bits you need for your server (including the operating system and additional software)
- Various configurations of CPU, memory, storage, and networking capacity for your instances, known as instance types
- Secure login information for your instances using key pairs (AWS stores the public key, and you store the private key in a secure place)
- Storage volumes for temporary data that's deleted when you stop or terminate your instance, known as instance store volumes
- Persistent storage volumes for your data using Amazon Elastic Block Store (Amazon EBS), known as Amazon EBS volumes
- Multiple physical locations for your resources, such as instances and Amazon EBS volumes, known as regions and Availability Zones
- A firewall that enables you to specify the protocols, ports, and source IP ranges that can reach your instances using security groups
- Static IP addresses for dynamic cloud computing, known as Elastic IP addresses
- Metadata, known as tags, that you can create and assign to your Amazon EC2 resources
- Virtual networks you can create that are logically isolated from the rest of the AWS cloud, and that you can optionally connect to your own network, known as virtual private clouds (VPCs).

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

**NEW QUESTION 190**

In AWS CloudHSM, in addition to the AWS recommendation that you use two or more HSM appliances in a high-availability configuration to prevent the loss of keys and data, you can also perform a remote backup/restore of a Luna SA partition if you have purchased a:

- A. Luna Restore HSM.
- B. Luna Backup HSM.

- C. Luna HSNI.
- D. Luna SA HSM.

**Answer:** B

**Explanation:**

In AWS CloudHSM, you can perform a remote backup/restore of a Luna SA partition if you have purchased a Luna Backup HSM.  
Reference: <http://docs.aws.amazon.com/cloudhsm/latest/userguide/cloud-hsm-backup-restore.html>

**NEW QUESTION 191**

AWS Identity and Access Management is a web service that enables Amazon Web Services (AWS) customers to manage users and user permissions in AWS. In addition to supporting IAM user policies, some services support resource-based permissions. Which of the following services are supported by resource-based permissions?

- A. Amazon SNS, and Amazon SQS and AWS Direct Connect.
- B. Amazon S3 and Amazon SQS and Amazon ElastiCache.
- C. Amazon S3, Amazon SNS, Amazon SQS, Amazon Glacier and Amazon EBS.
- D. Amazon Glacier, Amazon SNS, and Amazon CloudWatch

**Answer:** C

**Explanation:**

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, Amazon SQS, Amazon Glacier and Amazon EBS.  
Reference: [http://docs.aws.amazon.com/IAM/latest/UserGuide/Using\\_SpecificProducts.htm](http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_SpecificProducts.htm)

**NEW QUESTION 194**

A user is trying to launch a similar EC2 instance from an existing instance with the option "Launch More like this". The AMI of the selected instance is deleted. What will happen in this case?

- A. AWS does not need an AMI for the "Launch more like this" option
- B. AWS will launch the instance but will not create a new AMI
- C. AWS will create a new AMI and launch the instance
- D. AWS will throw an error saying that the AMI is deregistered

**Answer:** D

**Explanation:**

If the user has deregistered the AMI of an EC2 instance and is trying to launch a similar instance with the option "Launch more like this", AWS will throw an error saying that the AMI is deregistered or not available.  
Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/launching-instance.html>

**NEW QUESTION 197**

You are architecting a highly-scalable and reliable web application which will have a huge amount of content. You have decided to use CloudFront as you know it will speed up distribution of your static and dynamic web content and know that Amazon CloudFront integrates with Amazon CloudWatch metrics so that you can monitor your web application. Because you live in Sydney you have chosen the the Asia Pacific (Sydney) region in the AWS console. However you have set up this up but no CloudFront metrics seem to be appearing in the CloudWatch console. What is the most likely reason from the possible choices below for this?

- A. Metrics for CloudWatch are available only when you choose the same region as the application you are monitoring.
- B. You need to pay for CloudWatch for it to become active.
- C. Metrics for CloudWatch are available only when you choose the US East (Virginia)
- D. Metrics for CloudWatch are not available for the Asia Pacific region as yet

**Answer:** C

**Explanation:**

CloudFront is a global service, and metrics are available only when you choose the US East (N. Virginia) region in the AWS console. If you choose another region, no CloudFront metrics will appear in the CloudWatch console.

Reference:

<http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/monitoring-using-cloudwatch.html>

**NEW QUESTION 202**

A friend wants you to set up a small BitTorrent storage area for him on Amazon S3. You tell him it is highly unlikely that AWS would allow such a thing in their infrastructure. However you decide to investigate. Which of the following statements best describes using BitTorrent with Amazon S3?

- A. Amazon S3 does not support the BitTorrent protocol because it is used for pirated software.
- B. You can use the BitTorrent protocol but only for objects that are less than 100 GB in size.
- C. You can use the BitTorrent protocol but you need to ask AWS for specific permissions first.
- D. You can use the BitTorrent protocol but only for objects that are less than 5 GB in size

**Answer:** D

**Explanation:**

BitTorrent is an open, peer-to-peer protocol for distributing files. You can use the BitTorrent protocol to retrieve any publicly-accessible object in Amazon S3. Amazon S3 supports the BitTorrent protocol so that developers can save costs when distributing content at high scale. Amazon S3 is useful for simple, reliable storage of any data. The default distribution mechanism for Amazon S3 data is via client/server download. In client/server distribution, the entire object is transferred point-to-point from Amazon S3 to every authorized user who requests that object. While client/server delivery is appropriate for a wide variety of use

cases, it is not optimal for everybody. Specifically, the costs of client/server distribution increase linearly as the number of users downloading objects increases. This can make it expensive to distribute popular objects.

BitTorrent addresses this problem by recruiting the very clients that are downloading the object as distributors themselves: Each client downloads some pieces of the object from Amazon S3 and some from other clients, while simultaneously uploading pieces of the same object to other interested "peers." The benefit for publishers is that for large, popular files the amount of data actually supplied by Amazon S3 can be substantially lower than what it would have been serving the same clients via client/server download. Less data transferred means lower costs for the publisher of the object.

Reference: <http://docs.aws.amazon.com/AmazonS3/latest/dev/S3Torrent.html>

**NEW QUESTION 205**

A user has defined an AutoScaling termination policy to first delete the instance with the nearest billing hour. AutoScaling has launched 3 instances in the US-East-1A region and 2 instances in the US-East-1 B region. One of the instances in the US-East-1B region is running nearest to the billing hour. Which instance will AutoScaling terminate first while executing the termination action?

- A. Random Instance from US-East-1A
- B. Instance with the nearest billing hour in US-East-1 B
- C. Instance with the nearest billing hour in US-East-1A
- D. Random instance from US-East-1B

**Answer:** C

**Explanation:**

Even though the user has configured the termination policy, before AutoScaling selects an instance to terminate, it first identifies the Availability Zone that has more instances than the other Availability Zones used by the group. Within the selected Availability Zone, it identifies the instance that matches the specified termination policy.

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

**NEW QUESTION 206**

A user is planning a highly available application deployment with EC2. Which of the below mentioned options will not help to achieve HA?

- A. Elastic IP address
- B. PIOPS
- C. AMI
- D. Availability Zones

**Answer:** B

**Explanation:**

In Amazon Web Service, the user can achieve HA by deploying instances in multiple zones. The elastic IP helps the user achieve HA when one of the instances is down but still keeps the same URL. The AMI helps launching the new instance. The PIOPS is for the performance of EBS and does not help for HA. Reference: [http://media.amazonwebservices.com/AWS\\_Web\\_Hosting\\_Best\\_Practices.pdf](http://media.amazonwebservices.com/AWS_Web_Hosting_Best_Practices.pdf)

**NEW QUESTION 209**

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 210**

You have been given a scope to set up an AWS Media Sharing Framework for a new start up photo sharing company similar to flickr. The first thing that comes to mind about this is that it will obviously need a huge amount of persistent data storage for this framework. Which of the following storage options would be appropriate for persistent storage?

- A. Amazon Glacier or Amazon S3
- B. Amazon Glacier or AWS Import/Export
- C. AWS Import/Export or Amazon CloudFront
- D. Amazon EBS volumes or Amazon S3

**Answer:** D

**Explanation:**

Persistent storage-If you need persistent virtual disk storage similar to a physical disk drive for files or other data that must persist longer than the lifetime of a single Amazon EC2 instance, Amazon EBS volumes or Amazon S3 are more appropriate.

Reference: [http://media.amazonwebservices.com/AWS\\_Storage\\_Options.pdf](http://media.amazonwebservices.com/AWS_Storage_Options.pdf)

**NEW QUESTION 211**

After deploying a new website for a client on AWS, he asks if you can set it up so that if it fails it can be automatically redirected to a backup website that he has stored on a dedicated server elsewhere. You are wondering whether Amazon Route 53 can do this. Which statement below is correct in regards to Amazon Route 53?

- A. Amazon Route 53 can't help detect an outage
- B. You need to use another service.
- C. Amazon Route 53 can help detect an outage of your website and redirect your end users to alternate locations.
- D. Amazon Route 53 can help detect an outage of your website but can't redirect your end users to alternate locations.
- E. Amazon Route 53 can't help detect an outage of your website, but can redirect your end users to alternate locations.

**Answer: B**

**Explanation:**

With DNS Failover, Amazon Route 53 can help detect an outage of your website and redirect your end users to alternate locations where your application is operating properly.

Reference:

<http://aws.amazon.com/about-aws/whats-new/2013/02/11/announcing-dns-failover-for-route-53/>

**NEW QUESTION 216**

In Route 53, what does a Hosted Zone refer to?

- A. A hosted zone is a collection of geographical load balancing rules for Route 53.
- B. A hosted zone is a collection of resource record sets hosted by Route 53.
- C. A hosted zone is a selection of specific resource record sets hosted by CloudFront for distribution to Route 53.
- D. A hosted zone is the Edge Location that hosts the Route 53 records for a use

**Answer: B**

**Explanation:**

A Hosted Zone refers to a selection of resource record sets hosted by Route 53.

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

**NEW QUESTION 217**

Which of the following statements is true of Amazon EC2 security groups?

- A. You can change the outbound rules for EC2-Classi
- B. Also, you can add and remove rules to a group at any time.
- C. You can modify an existing rule in a grou
- D. However, you can't add and remove rules to a group.
- E. None of the statements are correct.
- F. You can't change the outbound rules for EC2-Classi
- G. However, you can add and remove rules to a group at any tim

**Answer: D**

**Explanation:**

When dealing with security groups, bear in mind that you can freely add and remove rules from a group, but you can't change the outbound rules for EC2-Classic. If you're using the Amazon EC2 console, you can modify existing rules, and you can copy the rules from an existing security group to a new security group.

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

**NEW QUESTION 218**

You need to create a management network using network interfaces for a virtual private cloud (VPC) network. Which of the following statements is incorrect pertaining to Best Practices for Configuring Network Interfaces.

- A. You can detach secondary (ethN) network interfaces when the instance is running or stoppe
- B. However, you can't detach the primary (eth0) interface.
- C. Launching an instance with multiple network interfaces automatically configures interfaces, private IP addresses, and route tables on the operating system of the instance.
- D. You can attach a network interface in one subnet to an instance in another subnet in the same VPC, however, both the network interface and the instance must reside in the same Availability Zone.
- E. Attaching another network interface to an instance is a valid method to increase or double the network bandwidth to or from the dual-homed instance

**Answer: D**

**Explanation:**

Best Practices for Configuring Network Interfaces

You can attach a network interface to an instance when it's running (hot attach), when it's stopped (warm attach), or when the instance is being launched (cold attach).

You can detach secondary (ethN) network interfaces when the instance is running or stopped. However, you can't detach the primary (eth0) interface.

You can attach a network interface in one subnet to an instance in another subnet in the same VPC, however, both the network interface and the instance must reside in the same Availability Zone.

When launching an instance from the CLI or API, you can specify the network interfaces to attach to the instance for both the primary (eth0) and additional network interfaces.

Launching an instance with multiple network interfaces automatically configures interfaces, private IP addresses, and route tables on the operating system of the instance.

A warm or hot attach of an additional network interface may require you to manually bring up the second interface, configure the private IP address, and modify the route table accordingly. (Instances running Amazon Linux automatically recognize the warm or hot attach and configure themselves.)

Attaching another network interface to an instance is not a method to increase or double the network bandwidth to or from the dual-homed instance.

Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html#use-network-and-security-apliances-in-your-vpc>

**NEW QUESTION 219**

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-Classic 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 224**

A user has launched an EC2 instance. The instance got terminated as soon as it was launched. Which of the below mentioned options is not a possible reason for this?

- A. The user account has reached the maximum volume limit
- B. The AM is missin
- C. It is the required part
- D. The snapshot is corrupt
- E. The user account has reached the maximum EC2 instance limit

**Answer: D**

**Explanation:**

When the user account has reached the maximum number of EC2 instances, it will not be allowed to launch an instance. AWS will throw an 'Instance Limit Exceeded' error. For all other reasons, such as

"AMI is missing part", "Corrupt Snapshot" or "Volume limit has reached" it will launch an EC2 instance and then terminate it.

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 227**

Can you encrypt EBS volumes?

- A. Yes, you can enable encryption when you create a new EBS volume using the AWS Management Console, API, or CLI.
- B. No, you should use a third-party software to perform raw block-level encryption of an EBS volume.
- C. Yes, but you must use a third-party API for encrypting data before it's loaded on EBS.
- D. Yes, you can encrypt with the special "ebs\_encrypt" command through Amazon API

**Answer: A**

**Explanation:**

With Amazon EBS encryption, you can now create an encrypted EBS volume and attach it to a supported instance type. Data on the volume, disk I/O, and snapshots created from the volume are then all encrypted. The encryption occurs on the servers that host the EC2 instances, providing encryption of data as it moves between EC2 instances and EBS storage. EBS encryption is based on the industry standard AES-256 cryptographic algorithm.

To get started, simply enable encryption when you create a new EBS volume using the AWS Management Console, API, or CLI. Amazon EBS encryption is available for all the latest EC2 instances in all commercially available AWS regions.

Reference:

<https://aws.amazon.com/about-aws/whats-new/2014/05/21/Amazon-EBS-encryption-now-available/>

**NEW QUESTION 228**

A user is making a scalable web application with compartmentalization. The user wants the log module to be able to be accessed by all the application functionalities in an asynchronous way. Each module of the application sends data to the log module, and based on the resource availability it will process the logs. Which AWS service helps this functionality?

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

**Answer: A**

**Explanation:**

Amazon Simple Queue Service (SQS) is a highly reliable distributed messaging system for storing messages as they travel between computers. By using Amazon SQS, developers can simply move data between distributed application components. It is used to achieve compartmentalization or loose coupling. In this case all the modules will send a message to the logger queue and the data will be processed by queue as per the resource availability.

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

**NEW QUESTION 232**

You have some very sensitive data stored on AWS S3 and want to try every possible alternative to keeping it secure in regards to access control. What are the mechanisms available for access control on AWS S3?

- A. (IAM) policies, Access Control Lists (ACLs), bucket policies, and query string authentication.
- B. (IAM) policies, Access Control Lists (ACLs) and bucket policies.
- C. Access Control Lists (ACLs), bucket policies, and query string authentication
- D. (IAM) policies, Access Control Lists (ACLs), bucket policies, query string authentication and encryption.

**Answer:** A

**Explanation:**

Amazon S3 supports several mechanisms that give you flexibility to control who can access your data as well as how, when, and where they can access it.

Amazon S3 provides four different access control mechanisms:

AWS Identity and Access Management (IAM) policies, Access Control Lists (ACLs), bucket policies, and query string authentication.

IAM enables organizations to create and manage multiple users under a single AWS account. With IAM policies, you can grant IAM users fine-grained control to your Amazon S3 bucket or objects. You can use ACLs to selectively add (grant) certain permissions on individual objects.

Amazon S3 bucket policies can be used to add or deny permissions across some or all of the objects within a single bucket.

With Query string authentication, you have the ability to share Amazon S3 objects through URLs that are valid for a specified period of time.

**NEW QUESTION 233**

Which one of the following can't be used as an origin server with Amazon CloudFront?

- A. A web server running in your infrastructure
- B. Amazon S3
- C. Amazon Glacier
- D. A web server running on Amazon EC2 instances

**Answer:** C

**Explanation:**

Amazon CloudFront is designed to work with Amazon S3 as your origin server, customers can also use Amazon CloudFront with origin servers running on Amazon EC2 instances or with any other custom origin.

Reference: <http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/distribution-web.html>

**NEW QUESTION 235**

You have written a CloudFormation template that creates 1 Elastic Load Balancer fronting 2 EC2 Instances. Which section of the template should you edit so that the DNS of the load balancer is returned upon creation of the stack?

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

**Answer:** B

**Explanation:**

You can use AWS CloudFormation's sample templates or create your own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run your application.

Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/outputs-section-structure.html>

**NEW QUESTION 238**

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 243**

You need to quickly set up an email-sending service because a client needs to start using it in the next hour. Amazon Simple Email Service (Amazon SES) seems to be the logical choice but there are several options available to set it up. Which of the following options to set up SES would best meet the needs of the client?

- A. Amazon SES console
- B. AWS CloudFormation
- C. SMTP Interface
- D. AWS Elastic Beanstalk

**Answer:** A

**Explanation:**

Amazon SES is an outbound-only email-sending service that provides an easy, cost-effective way for you to send email.

There are several ways that you can send an email by using Amazon SES. You can use the Amazon SES console, the Simple Mail Transfer Protocol (SMTP) interface, or you can call the Amazon SES API. Amazon SES console—This method is the quickest way to set up your system

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

**NEW QUESTION 248**

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 251**

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 252**

After deciding that EMR will be useful in analysing vast amounts of data for a gaming website that you are architecting you have just deployed an Amazon EMR Cluster and wish to monitor the cluster performance. Which of the following tools cannot be used to monitor the cluster performance?

- A. Kinesis
- B. Ganglia
- C. CloudWatch Metrics
- D. Hadoop Web Interfaces

**Answer:** A

**Explanation:**

Amazon EMR provides several tools to monitor the performance of your cluster. Hadoop Web Interfaces

Every cluster publishes a set of web interfaces on the master node that contain information about the cluster. You can access these web pages by using an SSH tunnel to connect them on the master node. For more information, see [View Web Interfaces Hosted on Amazon EMR Clusters](#).

CloudWatch Metrics

Every cluster reports metrics to CloudWatch. CloudWatch is a web service that tracks metrics, and which you can use to set alarms on those metrics. For more information, see [Monitor Metrics with CloudWatch](#).

Ganglia

Ganglia is a cluster monitoring tool. To have this available, you have to install Ganglia on the cluster when you launch it. After you've done so, you can monitor the cluster as it runs by using an SSH tunnel to connect to the Ganglia UI running on the master node. For more information, see [Monitor Performance with Ganglia](#).

Reference:

<http://docs.aws.amazon.com/ElasticMapReduce/latest/DeveloperGuide/emr-troubleshoot-tools.html>

**NEW QUESTION 253**

An application hosted at the EC2 instance receives an HTTP request from ELB. The same request has an X-Forwarded-For header, which has three IP addresses. Which system's IP will be a part of this header?

- A. Previous Request IP address.
- B. Client IP address.
- C. All of the answers listed here.
- D. Load Balancer IP address

**Answer:** C

**Explanation:**

When a user sends a request to ELB over HTTP/HTTPS, the request header log at the instance will only receive the IP of ELB. This is because ELB is the interceptor between the EC2 instance and the client request. To get the client IP, use the header X-Forwarded-For in header. The client IP address in the X-Forwarded-For request header is followed by the IP addresses of each successive proxy that passes along the request. The last IP address is the IP address that connects to the back-end application instance. e.g. if the HTTP request already has a header when it reaches the Load Balancer, the IP address from which the request came is appended at the end of the header followed by the IP address of the Load Balancer. In such cases, the X-Forwarded-For request header takes the following form:

X-Forwarded-For: clientIPAddress, previousRequestIPAddress, LoadBalancerIPAddress. Reference:

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

**NEW QUESTION 258**

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.html>

**NEW QUESTION 262**

A 3-tier e-commerce web application is currently deployed on-premises and will be migrated to AWS for greater scalability and elasticity. The web server currently shares read-only data using a network distributed file system. The app server tier uses a clustering mechanism for discovery and shared session state that depends on IP multicast. The database tier uses shared-storage clustering to provide database failover capability, and uses several read slaves for scaling. Data on all servers and the distributed file system directory is backed up weekly to off-site tapes. Which AWS storage and database architecture meets the requirements of the application?

- A. Web servers: store read-only data in EFS, and copy from EFS to root volume at boot time
- B. App servers: share state using a combination of DynamoDB and IP unicast
- C. Database: use RDS with multi-AZ deployment and one or more read replica
- D. Backup: web servers, app servers, and database backed up weekly to Glacier using snapshots.
- E. Web servers: store read-only data in an EC2 NFS server, mount to each web server at boot time
- F. App servers: share state using a combination of DynamoDB and IP multicast
- G. Database: use RDS with multi-AZ deployment and one or more Read Replica
- H. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.
- I. Web servers: store read-only data in EFS, and copy from EFS to root volume at boot time
- J. App servers: share state using a combination of DynamoDB and IP unicast
- K. Database: use RDS with multi-AZ deployment and one or more Read Replica
- L. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.
- M. Web servers: store read-only data in EFS, and copy from EFS to root volume at boot time
- N. App servers: share state using a combination of DynamoDB and IP unicast
- O. Database: use RDS with multi-AZ deployment
- P. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.

**Answer: C**

**Explanation:**

Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure (for example, instance hardware failure, storage failure, or network disruption), Amazon RDS performs an automatic failover to the standby, so that you can resume database operations as soon as the failover is complete. Since the endpoint for your DB Instance remains the same after a failover, your application can resume database operation without the need for manual administrative intervention.

**Benefits**

**Enhanced Durability**

Multi-AZ deployments for the MySQL, Oracle, and PostgreSQL engines utilize synchronous physical replication to keep data on the standby up-to-date with the primary. Multi-AZ deployments for the SQL Server engine use synchronous logical replication to achieve the same result, employing SQL Server-native Mirroring technology. Both approaches safeguard your data in the event of a DB Instance failure or loss of an Availability Zone.

If a storage volume on your primary fails in a Multi-AZ deployment, Amazon RDS automatically initiates a failover to the up-to-date standby. Compare this to a Single-AZ deployment: in case of a Single-AZ database failure, a user-initiated point-in-time-restore operation will be required. This operation can take several hours to complete, and any data updates that occurred after the latest restorable time (typically within the last five minutes) will not be available.

Amazon Aurora employs a highly durable, SSD-backed virtualized storage layer purpose-built for database workloads. Amazon Aurora automatically replicates your volume six ways, across three Availability Zones. Amazon Aurora storage is fault-tolerant, transparently handling the loss of up to two copies of data without affecting database write availability and up to three copies without affecting read availability. Amazon Aurora storage is also self-healing. Data blocks and disks are continuously scanned for errors and replaced automatically.

**Increased Availability**

You also benefit from enhanced database availability when running Multi-AZ deployments. If an Availability Zone failure or DB Instance failure occurs, your availability impact is limited to the time automatic failover takes to complete: typically under one minute for Amazon Aurora and one to two minutes for other database engines (see the RDS FAQ for details).

The availability benefits of Multi-AZ deployments also extend to planned maintenance and backups.

In the case of system upgrades like OS patching or DB Instance scaling, these operations are applied first on the standby, prior to the automatic failover. As a result, your availability impact is, again, only the time required for automatic failover to complete.

Unlike Single-AZ deployments, I/O activity is not suspended on your primary during backup for Multi-AZ deployments for the MySQL, Oracle, and PostgreSQL engines, because the backup is taken from the standby. However, note that you may still experience elevated latencies for a few minutes during backups for Multi-AZ deployments.

On instance failure in Amazon Aurora deployments, Amazon RDS uses RDS Multi-AZ technology to automate failover to one of up to 15 Amazon Aurora Replicas you have created in any of three Availability Zones. If no Amazon Aurora Replicas have been provisioned, in the case of a failure, Amazon RDS will attempt to create a new Amazon Aurora DB instance for you automatically.

**No Administrative Intervention**

DB Instance failover is fully automatic and requires no administrative intervention. Amazon RDS monitors the health of your primary and standbys, and initiates a failover automatically in response to a variety of failure conditions.

**Failover conditions**

Amazon RDS detects and automatically recovers from the most common failure scenarios for Multi-AZ deployments so that you can resume database operations as quickly as possible without administrative intervention. Amazon RDS automatically performs a failover in the event of any of the following:

- Loss of availability in primary Availability Zone
- Loss of network connectivity to primary Compute unit failure on primary
- Storage failure on primary

Note: When operations such as DB Instance scaling or system upgrades like OS patching are initiated for Multi-AZ deployments, for enhanced availability, they are applied first on the standby prior to an automatic failover. As a result, your availability impact is limited only to the time required for automatic failover to complete. Note that Amazon RDS Multi-AZ deployments do not failover automatically in response to database operations such as long running queries, deadlocks or database corruption errors.

**NEW QUESTION 265**

Your company has HQ in Tokyo and branch offices all over the world and is using a logistics software with a multi-regional deployment on AWS in Japan, Europe and USA. The logistic software has a 3-tier architecture and currently uses MySQL 5.6 for data persistence. Each region has deployed its own database. In the HQ region you run an hourly batch process reading data from every region to compute cross regional reports that are sent by email to all offices. This batch process must be completed as fast as possible to quickly optimize logistics. How do you build the database architecture in order to meet the requirements?

- A. For each regional deployment, use RDS MySQL with a master in the region and a read replica in the HQ region
- B. For each regional deployment, use MySQL on EC2 with a master in the region and send hourly EBS snapshots to the HQ region
- C. For each regional deployment, use RDS MySQL with a master in the region and send hourly RDS snapshots to the HQ region
- D. For each regional deployment, use MySQL on EC2 with a master in the region and use S3 to copy data files hourly to the HQ region
- E. Use Direct Connect to connect all regional MySQL deployments to the HQ region and reduce network latency for the batch process

**Answer: A**

**NEW QUESTION 270**

Your company is in the process of developing a next generation pet collar that collects biometric information to assist families with promoting healthy lifestyles for their pets. Each collar will push 30kb of biometric data in JSON format every 2 seconds to a collection platform that will process and analyze the data providing health trending information back to the pet owners and veterinarians via a web portal. Management has tasked you to architect the collection platform ensuring the following requirements are met.

Provide the ability for real-time analytics of the inbound biometric data. Ensure processing of the biometric data is highly durable, elastic and parallel. The results of the analytic processing should be persisted for data mining.

Which architecture outlined below will meet the initial requirements for the collection platform?

- A. Utilize S3 to collect the inbound sensor data, analyze the data from S3 with a daily scheduled Data Pipeline and save the results to a Redshift Cluster.
- B. Utilize Amazon Kinesis to collect the inbound sensor data, analyze the data with Kinesis clients and save the results to a Redshift cluster using EMR.
- C. Utilize SQS to collect the inbound sensor data, analyze the data from SQS with Amazon Kinesis and save the results to a Microsoft SQL Server RDS instance.
- D. Utilize EMR to collect the inbound sensor data, analyze the data from S3 with Amazon Kinesis and save the results to DynamoDB.

**Answer: B**

**NEW QUESTION 272**

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. But 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 274**

.....

## 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>