

Amazon

Exam Questions AWS-Solution-Architect-Associate

AWS Certified Solutions Architect - Associate



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

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 actMty 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

NEW QUESTION 3

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 4

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 5

A user is storing a large number of objects on AWS S3. The user wants to implement the search functionality among the objects. How can the user achieve this?

- A. Use the indexing feature of S3.
- B. Tag the objects with the metadata to search on that.
- C. Use the query functionality of S3.
- D. Make your own DB system which stores the S3 metadata for the search functionalit

Answer: D

Explanation: In Amazon Web Services, AWS S3 does not provide any query facility. To retrieve a specific object the user needs to know the exact bucket / object key. In this case it is recommended to have an own DB system which manages the S3 metadata and key mapping.
Reference: http://media.amazonwebservices.com/AWS_Storage_Options.pdf

NEW QUESTION 6

After setting up a Virtual Private Cloud (VPC) network, a more experienced cloud engineer suggests that to achieve low network latency and high network throughput you should look into setting up a placement group. You know nothing about this, but begin to do some research about it and are especially curious about its limitations. Which of the below statements is wrong in describing the limitations of a placement group?

- A. Although launching multiple instance types into a placement group is possible, this reduces the likelihood that the required capacity will be available for your launch to succeed.
- B. A placement group can span multiple Availability Zones.
- C. You can't move an existing instance into a placement group.
- D. A placement group can span peered VPCs

Answer: B

Explanation: A placement group is a logical grouping of 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. To provide the lowest latency, and the highest packet-per-second network performance for your placement group, choose an instance type that supports enhanced networking.

Placement groups have the following limitations:

The name you specify for a placement group a name must be unique within your AWS account. A placement group can't span multiple Availability Zones.

Although launching multiple instance types into a placement group is possible, this reduces the likelihood that the required capacity will be available for your launch to succeed. We recommend using the same instance type for all instances in a placement group.

You can't merge placement groups. Instead, you must terminate the instances in one placement group, and then relaunch those instances into the other placement group.

A placement group can span peered VPCs; however, you will not get full-bisection bandwidth between instances in peered VPCs. For more information about VPC peering connections, see VPC Peering in the Amazon VPC User Guide.

You can't move an existing instance into a placement group. You can create an AM from your existing instance, and then launch a new instance from the AMI into a placement group.

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

NEW QUESTION 7

You are migrating an internal server on your DC to an EC2 instance with EBS volume. Your server disk usage is around 500GB so you just copied all your data to a 2TB disk to be used with AWS Import/Export. Where will the data be imported once it arrives at Amazon?

- A. to a 2TB EBS volume
- B. to an S3 bucket with 2 objects of 1TB
- C. to an 500GB EBS volume
- D. to an S3 bucket as a 2TB snapshot

Answer: B

Explanation: An import to Amazon EBS will have different results depending on whether the capacity of your storage device is less than or equal to 1 TB or greater than 1 TB. The maximum size of an Amazon EBS snapshot is 1 TB, so if the device image is larger than 1 TB, the image is chunked and stored on Amazon S3. The target location is determined based on the total capacity of the device, not the amount of data on the device.

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

NEW QUESTION 8

True or false? A VPC contains multiple subnets, where each subnet can span multiple Availability Zones.

- A. This is true only if requested during the set-up of VPC.
- B. This is true.
- C. This is false.
- D. This is true only for US region

Answer: C

Explanation: A VPC can span several Availability Zones. In contrast, a subnet must reside within a single Availability Zone.

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

NEW QUESTION 9

You are looking at ways to improve some existing infrastructure as it seems a lot of engineering resources are being taken up with basic management and monitoring tasks and the costs seem to be excessive.

You are thinking of deploying Amazon ElastiCache to help. Which of the following statements is true in regards to ElastiCache?

- A. You can improve load and response times to user actions and queries however the cost associated with scaling web applications will be more.
- B. You can't improve load and response times to user actions and queries but you can reduce the cost associated with scaling web applications.
- C. You can improve load and response times to user actions and queries however the cost associated with scaling web applications will remain the same.
- D. You can improve load and response times to user actions and queries and also reduce the cost associated with scaling web applications.

Answer: D

Explanation: Amazon ElastiCache is a web service that makes it easy to deploy and run Memcached or Redis protocol-compliant server nodes in the cloud.

Amazon ElastiCache improves the performance of web applications by allowing you to retrieve information from a fast, managed, in-memory caching system, instead of relying entirely on slower disk-based databases. The service simplifies and offloads the management, monitoring and operation of in-memory cache environments, enabling your engineering resources to focus on developing applications. Using Amazon ElastiCache, you can not only improve load and response times to user actions and queries, but also reduce the cost associated with scaling web applications.

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

NEW QUESTION 10

Your supervisor has asked you to build a simple file synchronization service for your department. He doesn't want to spend too much money and he wants to be notified of any changes to files by email. What do you think would be the best Amazon service to use for the email solution?

- A. Amazon SES
- B. Amazon CloudSearch
- C. Amazon SWF
- D. Amazon AppStream

Answer: A

Explanation: File change notifications can be sent via email to users following the resource with Amazon Simple Email Service (Amazon SES), an easy-to-use, cost-effective email solution.

Reference: http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_filesync_08.pdf

NEW QUESTION 10

Does DynamoDB support in-place atomic updates?

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

Answer: A

Explanation: DynamoDB supports in-place atomic updates.

Reference:

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

NEW QUESTION 11

Your manager has just given you access to multiple VPN connections that someone else has recently set up between all your company's offices. She needs you to make sure that the communication between the VPNs is secure. Which of the following services would be best for providing a low-cost hub-and-spoke model for primary or backup connectivity between these remote offices?

- A. Amazon CloudFront
- B. AWS Direct Connect
- C. AWS CloudHSM
- D. AWS VPN CloudHub

Answer: D

Explanation: If you have multiple VPN connections, you can provide secure communication between sites using the AWS VPN CloudHub. The VPN CloudHub operates on a simple hub-and-spoke model that you can use with or without a VPC. This design is suitable for customers with multiple branch offices and existing Internet connections who would like to implement a convenient, potentially low-cost hub-and-spoke model for primary or backup connectivity between these remote offices.

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

NEW QUESTION 15

Does Amazon DynamoDB support both increment and decrement atomic operations?

- A. Only increment, since decrement are inherently impossible with DynamoDB's data model.
- B. No, neither increment nor decrement operations.
- C. Yes, both increment and decrement operations.
- D. Only decrement, since increment are inherently impossible with DynamoDB's data model.

Answer: C

Explanation: Amazon DynamoDB supports increment and decrement atomic operations.

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

NEW QUESTION 20

An organization has three separate AWS accounts, one each for development, testing, and production. The organization wants the testing team to have access to certain AWS resources in the production account. How can the organization achieve this?

- A. It is not possible to access resources of one account with another account.
- B. Create the IAM roles with cross account access.

- C. Create the IAM user in a test account, and allow it access to the production environment with the IAM policy.
- D. Create the IAM users with cross account acces

Answer: B

Explanation: An organization has multiple AWS accounts to isolate a development environment from a testing or production environment. At times the users from one account need to access resources in the other account, such as promoting an update from the development environment to the production environment. In this case the IAM role with cross account access will provide a solution. Cross account access lets one account share access to their resources with users in the other AWS accounts.

Reference: http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf

NEW QUESTION 23

A user has created an EBS volume with 1000 IOPS. What is the average IOPS that the user will get for most of the year as per EC2 SLA if the instance is attached to the EBS optimized instance?

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

Answer: D

Explanation: As per AWS SLA if the instance is attached to an EBS-Optimized instance, then the Provisioned IOPS volumes are designed to deliver within 10% of the provisioned IOPS performance 99.9% of the time in a given year. Thus, if the user has created a volume of 1000 IOPS, the user will get a minimum 900 IOPS 99.9% time of the year.

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

NEW QUESTION 27

You are in the process of creating a Route 53 DNS failover to direct traffic to two EC2 zones. Obviously, if one fails, you would like Route 53 to direct traffic to the other region. Each region has an ELB with some instances being distributed. What is the best way for you to configure the Route 53 health check?

- A. Route 53 doesn't support ELB with an internal health check. You need to create your own Route 53 health check of the ELB
- B. Route 53 natively supports ELB with an internal health check
- C. Turn "Evaluate target health" off and "Associate with Health Check" on and Route 53 will use the ELB's internal health check.
- D. Route 53 doesn't support ELB with an internal health check
- E. You need to associate your resource record set for the ELB with your own health check
- F. Route 53 natively supports ELB with an internal health check
- G. Turn "Evaluate target health" on and "Associate with Health Check" off and Route 53 will use the ELB's internal health check.

Answer: D

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. When you enable this feature, Route 53 uses health checks-regularly making Internet requests to your application's endpoints from multiple locations around the world-to determine whether each endpoint of your application is up or down.

To enable DNS Failover for an ELB endpoint, create an Alias record pointing to the ELB and set the "Evaluate Target Health" parameter to true. Route 53 creates and manages the health checks for your ELB automatically. You do not need to create your own Route 53 health check of the ELB. You also do not need to associate your resource record set for the ELB with your own health check, because Route 53 automatically associates it with the health checks that Route 53 manages on your behalf. The ELB health check will also inherit the health of your backend instances behind that ELB.

Reference:

<http://aws.amazon.com/about-aws/whats-new/2013/05/30/amazon-route-53-adds-elb-integration-for-dns-failover/>

NEW QUESTION 32

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 33

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 time

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 34

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 37

In Amazon EC2, partial instance-hours are billed .

- A. per second used in the hour
- B. per minute used
- C. by combining partial segments into full hours
- D. as full hours

Answer: D

Explanation: Partial instance-hours are billed to the next hour. Reference: <http://aws.amazon.com/ec2/faqs/>

NEW QUESTION 38

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 lost

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 41

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 subnet

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

NEW QUESTION 42

While using the EC2 GET requests as URLs, the is the URL that serves as the entry point for the web service.

- A. token
- B. endpoint
- C. action
- D. None of these

Answer: B

Explanation: The endpoint is the URL that serves as the entry point for the web service.
Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-query-api.html>

NEW QUESTION 43

You have been asked to build a database warehouse using Amazon Redshift. You know a little about it, including that it is a SQL data warehouse solution, and uses industry standard ODBC and JDBC connections and PostgreSQL drivers. However you are not sure about what sort of storage it uses for database tables. What sort of storage does Amazon Redshift use for database tables?

- A. InnoDB Tables
- B. NDB data storage
- C. Columnar data storage
- D. NDB CLUSTER Storage

Answer: C

Explanation: Amazon Redshift achieves efficient storage and optimum query performance through a combination of massively parallel processing, columnar data storage, and very efficient, targeted data compression encoding schemes. Columnar storage for database tables is an important factor in optimizing analytic query performance because it drastically reduces the overall disk I/O requirements and reduces the amount of data you need to load from disk.
Reference: http://docs.aws.amazon.com/redshift/latest/dg/c_columnar_storage_disk_mem_mgmt.html

NEW QUESTION 47

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 50

Which of the below mentioned options is not available when an instance is launched by Auto Scaling with EC2 Classic?

- A. Public IP
- B. Elastic IP
- C. Private DNS
- D. Private IP

Answer: B

Explanation: Auto Scaling supports both EC2 classic and EC2-VPC. When an instance is launched as a part of EC2 classic, it will have the public IP and DNS as well as the private IP and DNS.
Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/GettingStartedTutorial.html>

NEW QUESTION 53

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 55

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

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

Answer: D

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

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

NEW QUESTION 58

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 99.99999999% 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 59

Your EBS volumes do not seem to be performing as expected and your team leader has requested you look into improving their performance. Which of the following is not a true statement relating to the performance of your EBS volumes?

- A. Frequent snapshots provide a higher level of data durability and they will not degrade the performance of your application while the snapshot is in progress.
- B. General Purpose (SSD) and Provisioned IOPS (SSD) volumes have a throughput limit of 128 MB/s per volume.
- C. There is a relationship between the maximum performance of your EBS volumes, the amount of I/O you are drMng to them, and the amount of time it takes for each transaction to complete.
- D. There is a 5 to 50 percent reduction in IOPS when you first access each block of data on a newly created or restored EBS volume

Answer: A

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

Frequent snapshots provide a higher level of data durability, but they may slightly degrade the performance of your application while the snapshot is in progress. This trade off becomes critical when you have data that changes rapidly. Whenever possible, plan for snapshots to occur during off-peak times in order to minimize workload impact.

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

NEW QUESTION 64

A major finance organisation has engaged your company to set up a large data mining application. Using AWS you decide the best service for this is Amazon Elastic MapReduce(EMR) which you know uses Hadoop. Which of the following statements best describes Hadoop?

- A. Hadoop is 3rd Party software which can be installed using AMI
- B. Hadoop is an open source python web framework
- C. Hadoop is an open source Java software framework
- D. Hadoop is an open source javascript framework

Answer: C

Explanation: Amazon EMR uses Apache Hadoop as its distributed data processing engine.

Hadoop is an open source, Java software framework that supports data-intensive distributed applications running on large clusters of commodity hardware.

Hadoop implements a programming model named "MapReduce," where the data is dMded into many small fragments of work, each of which may be executed on any node in the cluster.

This framework has been widely used by developers, enterprises and startups and has proven to be a reliable software platform for processing up to petabytes of data on clusters of thousands of commodity machines.

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

NEW QUESTION 69

In Amazon EC2 Container Service, are other container types supported?

- A. Yes, EC2 Container Service supports any container service you need.
- B. Yes, EC2 Container Service also supports Microsoft container service.
- C. No, Docker is the only container platform supported by EC2 Container Service presently.
- D. Yes, EC2 Container Service supports Microsoft container service and Openstac

Answer: C

Explanation: In Amazon EC2 Container Service, Docker is the only container platform supported by EC2 Container Service presently.

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

NEW QUESTION 74

As AWS grows, most of your clients' main concerns seem to be about security, especially when all of their competitors also seem to be using AWS. One of your clients asks you whether having a competitor who hosts their EC2 instances on the same physical host would make it easier for the competitor to hack into the client's data. Which of the following statements would be the best choice to put your client's mind at rest?

- A. Different instances running on the same physical machine are isolated from each other via a 256-bit Advanced Encryption Standard (AES-256).
- B. Different instances running on the same physical machine are isolated from each other via the Xen hypervisor and via a 256-bit Advanced Encryption Standard (AES-256).
- C. Different instances running on the same physical machine are isolated from each other via the Xen hypervisor.
- D. Different instances running on the same physical machine are isolated from each other via IAM permissions.

Answer: C

Explanation: Amazon Elastic Compute Cloud (EC2) is a key component in Amazon's Infrastructure as a Service (IaaS), providing resizable computing capacity using server instances in AWS's data centers. Amazon EC2 is designed to make web-scale computing easier by enabling you to obtain and configure capacity with minimal friction.

You create and launch instances, which are collections of platform hardware and software. Different instances running on the same physical machine are isolated from each other via the Xen hypervisor.

Amazon is active in the Xen community, which provides awareness of the latest developments. In addition, the AWS firewall resides within the hypervisor layer, between the physical network interface and the instance's virtual interface. All packets must pass through this layer, thus an instance's neighbors have no more access to that instance than any other host on the Internet and can be treated as if they are on separate physical hosts. The physical RAM is separated using similar mechanisms.

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

NEW QUESTION 79

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 83

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. Ye
- B. Amazon SimpleDB
- C. N
- D. Not presently
- E. Ye
- F. Amazon Redshift
- G. Ye
- H. Your choice of relational AMLs 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

NEW QUESTION 84

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 85

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 90

Can resource record sets in a hosted zone have a different domain suffix (for example, www.blog.acme.com and www.acme.ca)?

- A. Yes, it can have for a maximum of three different TLDs.
- B. Yes
- C. Yes, it can have depending on the TLD.
- D. No

Answer: D

Explanation: The resource record sets contained in a hosted zone must share the same suffix. For example, the example.com hosted zone can contain resource record sets for www.example.com and www.aws.example.com subdomains, but it cannot contain resource record sets for a www.example.ca subdomain.
Reference: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/AboutHostedZones.html>

NEW QUESTION 95

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

NEW QUESTION 98

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 103

Which of the following AWS CLI commands is syntactically incorrect?

1. `$ aws ec2 describe-instances`
2. `$ aws ec2 start-instances --instance-ids i-1348636c`
3. `$ aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError -message "Script Failure"`
4. `$ aws sqs receive-message --queue-url https://queue.amazonaws.com/546419318123/Test`

- A. 3
- B. 4
- C. 2
- D. 1

Answer: A

Explanation: The following CLI command is missing a hyphen before "-message".

```
aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError -message "Script Failure"
```

It has been added below in red

```
aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError ---message "Script Failure"
```

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

NEW QUESTION 104

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 106

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 107

An online gaming site asked you if you can deploy a database that is a fast, highly scalable NoSQL database service in AWS for a new site that he wants to build. Which database should you recommend?

- A. Amazon DynamoDB
- B. Amazon RDS
- C. Amazon Redshift
- D. Amazon SimpleDB

Answer: A

Explanation: Amazon DynamoDB is ideal for database applications that require very low latency and predictable performance at any scale but don't need complex querying capabilities like joins or transactions. Amazon DynamoDB is a fully-managed NoSQL database service that offers high performance, predictable throughput and low cost. It is easy to set up, operate, and scale.

With Amazon DynamoDB, you can start small, specify the throughput and storage you need, and easily scale your capacity requirements on the fly. Amazon DynamoDB automatically partitions data over a number of servers to meet your request capacity. In addition, DynamoDB automatically replicates your data synchronously across multiple Availability Zones within an AWS Region to ensure high-availability and data durability.

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

NEW QUESTION 110

You have been doing a lot of testing of your VPC Network by deliberately failing EC2 instances to test whether instances are failing over properly. Your customer who will be paying the AWS bill for all this asks you if he being charged for all these instances. You try to explain to him how the billing works on EC2 instances to the best of your knowledge. What would be an appropriate response to give to the customer in regards to this?

- A. Billing commences when Amazon EC2 AM instance is completely up and billing ends as soon as the instance starts to shutdown.
- B. Billing only commences only after 1 hour of uptime and billing ends when the instance terminates.
- C. Billing commences when Amazon EC2 initiates the boot sequence of an AM instance and billing ends when the instance shuts down.
- D. Billing commences when Amazon EC2 initiates the boot sequence of an AM instance and billing ends as soon as the instance starts to shutdown.

Answer: C

Explanation: Billing commences when Amazon EC2 initiates the boot sequence of an AM instance. Billing ends when the instance shuts down, which could occur through a web services command, by running "shutdown -h", or through instance failure.

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

NEW QUESTION 112

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

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

Answer: D

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

Reference:

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

NEW QUESTION 113

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 116

You need to change some settings on Amazon Relational Database Service but you do not want the database to reboot immediately which you know might happen depending on the setting that you change. Which of the following will cause an immediate DB instance reboot to occur?

- A. You change storage type from standard to PIOPS, and Apply Immediately is set to true.
- B. You change the DB instance class, and Apply Immediately is set to false.
- C. You change a static parameter in a DB parameter group.
- D. You change the backup retention period for a DB instance from 0 to a nonzero value or from a nonzero value to 0, and Apply Immediately is set to false.

Answer: A

Explanation: A DB instance outage can occur when a DB instance is rebooted, when the DB instance is put into a state that prevents access to it, and when the database is restarted. A reboot can occur when you manually reboot your DB instance or when you change a DB instance setting that requires a reboot before it can take effect.

A DB instance reboot occurs immediately when one of the following occurs:

You change the backup retention period for a DB instance from 0 to a nonzero value or from a nonzero value to 0 and set Apply Immediately to true.

You change the DB instance class, and Apply Immediately is set to true.

You change storage type from standard to PIOPS, and Apply Immediately is set to true.

A DB instance reboot occurs during the maintenance window when one of the following occurs:

You change the backup retention period for a DB instance from 0 to a nonzero value or from a nonzero value to 0, and Apply Immediately is set to false.

You change the DB instance class, and Apply Immediately is set to false. Reference:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Troubleshooting.htm#CHAP_Troubleshooting.Security

NEW QUESTION 121

What does the following policy for Amazon EC2 do?

```
{
  "Statement": [{
    "Effect": "Allow", "Action": "ec2:Describe*", "Resource": "*"
  }]
}
```

- A. Allow users to use actions that start with "Describe" over all the EC2 resources.
- B. Share an AMI with a partner
- C. Share an AMI within the account
- D. Allow a group to only be able to describe, run, stop, start, and terminate instances

Answer: A

Explanation: You can use IAM policies to control the actions that your users can perform against your EC2 resources. For instance, a policy with the following statement will allow users to perform actions whose name start with "Describe" against all your EC2 resources.

```
{
  "Statement": [{
    "Effect": "Allow", "Action": "ec2:Describe*", "Resource": "*"
  }]
}
```

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

NEW QUESTION 126

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

NEW QUESTION 127

Amazon RDS provides high availability and failover support for DB instances using .

- A. customized deployments
- B. Appstream customizations
- C. log events

D. Multi-AZ deployments

Answer: D

Explanation: Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. Multi-AZ deployments for Oracle, PostgreSQL, MySQL, and MariaDB DB instances use Amazon technology, while SQL Server DB instances use SQL Server Mirroring.
Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.IV|ultiAZ.html>

NEW QUESTION 130

What does Amazon DynamoDB provide?

- A. A predictable and scalable MySQL database
- B. A fast and reliable PL/SQL database cluster
- C. A standalone Cassandra database, managed by Amazon Web Services
- D. A fast, highly scalable managed NoSQL database service

Answer: D

Explanation: Amazon DynamoDB is a managed NoSQL database service offered by Amazon. It automatically manages tasks like scalability for you while it provides high availability and durability for your data, allowing you to concentrate in other aspects of your application.
Reference: check link - https://aws.amazon.com/running_databases/

NEW QUESTION 135

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 137

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 142

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. What formatting is required for this template?

- A. JSON-formatted document
- B. CSS-formatted document
- C. XML-formatted document
- D. HTML-formatted document

Answer: A

Explanation: You can write an AWS CloudFormation template (a JSON-formatted document) in a text editor or pick an existing template. The template describes the resources you want and their settings. For example, suppose you want to create an Amazon EC2. Your template can declare an instance Amazon EC2 and describe its properties, as shown in the following example:

```
{
  "AWSTemplateFormatVersion" : "2010-09-09",
  "Description" : "A simple Amazon EC2 instance", "Resources" : {
  "MyEC2Instance" : {
    "Type" : "AWS::EC2::Instance", "Properties" : {
    "ImageId" : "ami-2f726546", "InstanceType" : "t1.micro"
    }
  }
}
}
```

Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-what-is-how-does-it-work.html>

NEW QUESTION 147

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

NEW QUESTION 149

You decide that you need to create a number of Auto Scaling groups to try and save some money as you have noticed that at certain times most of your EC2 instances are not being used. By default, what is the maximum number of Auto Scaling groups that AWS will allow you to create?

- A. 12
- B. Unlimited
- C. 20
- D. 2

Answer: C

Explanation: Auto Scaling is an AWS service that allows you to increase or decrease the number of EC2 instances within your application's architecture. With Auto Scaling, you create collections of EC2 instances, called Auto Scaling groups. You can create these groups from scratch, or from existing EC2 instances that are already in production.

Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_autoscaling

NEW QUESTION 153

A user needs to run a batch process which runs for 10 minutes. This will only be run once, or at maximum twice, in the next month, so the processes will be temporary only. The process needs 15 X-Large instances. The process downloads the code from S3 on each instance when it is launched, and then generates a temporary log file. Once the instance is terminated, all the data will be lost. Which of the below mentioned pricing models should the user choose in this case?

- A. Spot instance.
- B. Reserved instance.
- C. On-demand instance.
- D. EBS optimized instance

Answer: A

Explanation: In Amazon Web Services, the spot instance is useful when the user wants to run a process temporarily. The spot instance can terminate the instance if the other user outbids the existing bid. In this case all storage is temporary and the data is not required to be persistent. Thus, the spot instance is a good option to save money.

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

NEW QUESTION 157

You have been storing massive amounts of data on Amazon Glacier for the past 2 years and now start to wonder if there are any limitations on this. What is the correct answer to your QUESTION ?

- A. The total volume of data is limited but the number of archives you can store are unlimited.
- B. The total volume of data is unlimited but the number of archives you can store are limited.
- C. The total volume of data and number of archives you can store are unlimited.
- D. The total volume of data is limited and the number of archives you can store are limited.

Answer: C

Explanation: An archive is a durably stored block of information. You store your data in Amazon Glacier as archives. You may upload a single file as an archive, but your costs will be lower if you aggregate your data. TAR and ZIP are common formats that customers use to aggregate multiple files into a single file before uploading to Amazon Glacier.

The total volume of data and number of archives you can store are unlimited. Individual Amazon Glacier archives can range in size from 1 byte to 40 terabytes. The largest archive that can be uploaded in a single upload request is 4 gigabytes.

For items larger than 100 megabytes, customers should consider using the Multipart upload capability. Archives stored in Amazon Glacier are immutable, i.e. archives can be uploaded and deleted but cannot be edited or overwritten.
Reference: <https://aws.amazon.com/glacier/faqs/>

NEW QUESTION 158

You are setting up your first Amazon Virtual Private Cloud (Amazon VPC) so you decide to use the VPC wizard in the AWS console to help make it easier for you. Which of the following statements is correct regarding instances that you launch into a default subnet via the VPC wizard?

- A. Instances that you launch into a default subnet receive a public IP address and 10 private IP addresses.
- B. Instances that you launch into a default subnet receive both a public IP address and a private IP address.
- C. Instances that you launch into a default subnet don't receive any IP addresses and you need to define them manually.
- D. Instances that you launch into a default subnet receive a public IP address and 5 private IP addresses.

Answer: B

Explanation: Instances that you launch into a default subnet receive both a public IP address and a private IP address. Instances in a default subnet also receive both public and private DNS hostnames. Instances that you launch into a nondefault subnet in a default VPC don't receive a public IP address or a DNS hostname. You can change your subnet's default public IP addressing behavior.

Reference: <http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/default-vpc.html>

NEW QUESTION 159

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 161

Does AWS CloudFormation support Amazon EC2 tagging?

- A. Yes, AWS CloudFormation supports Amazon EC2 tagging
- B. No, CloudFormation doesn't support any tagging
- C. No, it doesn't support Amazon EC2 tagging.
- D. It depends if the Amazon EC2 tagging has been defined in the template

Answer: A

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

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

NEW QUESTION 162

An existing client comes to you and says that he has heard that launching instances into a VPC (virtual private cloud) is a better strategy than launching instances into a EC2-classic which he knows is what you currently do. You suspect that he is correct and he has asked you to do some research about this and get back to him. Which of the following statements is true in regards to what ability launching your instances into a VPC instead of EC2-Classic gives you?

- A. All of the things listed here.
- B. Change security group membership for your instances while they're running
- C. Assign static private IP addresses to your instances that persist across starts and stops
- D. Define network interfaces, and attach one or more network interfaces to your instances

Answer: A

Explanation: By launching your instances into a VPC instead of EC2-Classic, you gain the ability to: Assign static private IP addresses to your instances that persist across starts and stops Assign multiple IP addresses to your instances

Define network interfaces, and attach one or more network interfaces to your instances Change security group membership for your instances while they're running

Control the outbound traffic from your instances (egress filtering) in addition to controlling the inbound traffic to them (ingress filtering)

Add an additional layer of access control to your instances in the form of network access control lists (ACL)

Run your instances on single-tenant hardware

Reference: http://media.amazonwebservices.com/AWS_Cloud_Best_Practices.pdf

NEW QUESTION 166

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

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

Answer: B

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

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

NEW QUESTION 169

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 174

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

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

Answer: B

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

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

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

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

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

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

NEW QUESTION 175

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

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

Answer: C

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

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

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

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

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

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

NEW QUESTION 178

A user has created a subnet in VPC and launched an EC2 instance within it. The user has not selected the option to assign the IP address while launching the instance. The user has 3 elastic IPs and is trying to assign one of the Elastic IPs to the VPC instance from the console. The console does not show any instance in the IP assignment screen. What is a possible reason that the instance is unavailable in the assigned IP console?

- A. The IP address may be attached to one of the instances
- B. The IP address belongs to a different zone than the subnet zone
- C. The user has not created an internet gateway
- D. The IP addresses belong to EC2 Classic; so they cannot be assigned to VPC

Answer: D

Explanation: A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. When the user is launching an instance he needs to select an option which attaches a public IP to the instance. If the user has not selected the option to attach the public IP then it will only have a private IP when launched. If the user wants to connect to an instance from the internet he should create an elastic IP with VPC. If the elastic IP is a part of EC2 Classic it cannot be assigned to a VPC instance.

Reference: <http://docs.aws.amazon.com/AmazonVPC/latest/GettingStartedGuide/LaunchInstance.html>

NEW QUESTION 179

An accountant asks you to design a small VPC network for him and, due to the nature of his business, just needs something where the workload on the network will be low, and dynamic data will be accessed infrequently. Being an accountant, low cost is also a major factor. Which EBS volume type would best suit his requirements?

- A. Magnetic
- B. Any, as they all perform the same and cost the same.
- C. General Purpose (SSD)
- D. Magnetic or Provisioned IOPS (SSD)

Answer: A

Explanation: You can choose between three EBS volume types to best meet the needs of their workloads: General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic. General Purpose (SSD) is the new, SSD-backed, general purpose EBS volume type that we recommend as the default choice for customers. General Purpose (SSD) volumes are suitable for a broad range of workloads, including small to medium sized databases, development and test environments, and boot volumes. Provisioned IOPS (SSD) volumes offer storage with consistent and low-latency performance, and are designed for I/O intensive applications such as large relational or NoSQL databases. Magnetic volumes provide the lowest cost per gigabyte of all EBS volume types. Magnetic volumes are ideal for workloads where data is accessed infrequently, and applications where the lowest storage cost is important.

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

NEW QUESTION 181

A user has launched one EC2 instance in the US East region and one in the US West region. The user has launched an RDS instance in the US East region. How can the user configure access from both the EC2 instances to RDS?

- A. It is not possible to access RDS of the US East region from the US West region
- B. Configure the US West region's security group to allow a request from the US East region's instance and configure the RDS security group's ingress rule for the US East EC2 group
- C. Configure the security group of the US East region to allow traffic from the US West region's instance and configure the RDS security group's ingress rule for the US East EC2 group
- D. Configure the security group of both instances in the ingress rule of the RDS security group

Answer: C

Explanation: The user cannot authorize an Amazon EC2 security group if it is in a different AWS Region than the RDS DB instance. The user can authorize an IP range or specify an Amazon EC2 security group in the same region that refers to an IP address in another region. In this case allow IP of US West inside US East's security group and open the RDS security group for US East region.

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

NEW QUESTION 186

In Amazon EC2, if your EBS volume stays in the detaching state, you can force the detachment by clicking .

- A. Force Detach
- B. Detach Instance
- C. AttachVolume
- D. AttachInstance

Answer: A

Explanation: If your volume stays in the detaching state, you can force the detachment by clicking Force Detach. Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/ebs-detaching-volume.html>

NEW QUESTION 189

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

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

Answer: A

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

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

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

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

NEW QUESTION 194

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 199

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 202

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

NEW QUESTION 206

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

NEW QUESTION 208

Select the correct statement: Within Amazon EC2, when using Linux instances, the device name /dev/sda1 is .

- A. reserved for EBS volumes
- B. recommended for EBS volumes
- C. recommended for instance store volumes
- D. reserved for the root device

Answer: D

Explanation: Within Amazon EC2, when using a Linux instance, the device name /dev/sda1 is reserved for the root device.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/device_naming.html

NEW QUESTION 211

The common use cases for DynamoDB Fine-Grained Access Control (FGAC) are cases in which the end user wants .

- A. to change the hash keys of the table directly
- B. to check if an IAM policy requires the hash keys of the tables directly
- C. to read or modify any codecommit key of the table directly, without a middle-tier service
- D. to read or modify the table directly, without a middle-tier service

Answer: D

Explanation: FGAC can benefit any application that tracks information in a DynamoDB table, where the end user (or application client acting on behalf of an end user) wants to read or modify the table directly, without a middle-tier service. For instance, a developer of a mobile app named Acme can use FGAC to track the top score of every Acme user in a DynamoDB table. FGAC allows the application client to modify only the top score for the user that is currently running the application.
Reference: http://aws.amazon.com/dynamodb/faqs/#security_anchor

NEW QUESTION 214

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 218

You need to set up security for your VPC and you know that Amazon VPC provides two features that you can use to increase security for your VPC: security groups and network access control lists (ACLs). You have already looked into security groups and you are now trying to understand ACLs. Which statement below is incorrect in relation to ACLs?

- A. Supports allow rules and deny rules.
- B. Is stateful: Return traffic is automatically allowed, regardless of any rules.
- C. Processes rules in number order when deciding whether to allow traffic.
- D. Operates at the subnet level (second layer of defense).

Answer: B

Explanation: Amazon VPC provides two features that you can use to increase security for your VPC:
Security groups—Act as a firewall for associated Amazon EC2 instances, controlling both inbound and outbound traffic at the instance level
Network access control lists (ACLs)—Act as a firewall for associated subnets, controlling both inbound and outbound traffic at the subnet level
Security groups are stateful: (Return traffic is automatically allowed, regardless of any rules) Network ACLs are stateless: (Return traffic must be explicitly allowed by rules)
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Security.html

NEW QUESTION 223

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

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

Answer: B

Explanation: Amazon Simple Notification Service (Amazon SNS) is a fast, flexible, and fully managed push messaging service. Amazon SNS can be used to make push notifications to mobile devices. Amazon SNS can deliver notifications by SMS text message or email to the Amazon Simple Queue Service (SQS) queues or to any HTTP endpoint. In this case user can use the SNS apis to send SMS.

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

NEW QUESTION 224

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 229

You have a lot of data stored in the AWS Storage Gateway and your manager has come to you asking about how the billing is calculated, specifically the Virtual Tape Shelf usage. What would be a correct response to this?

- A. You are billed for the virtual tape data you store in Amazon Glacier and are billed for the size of the virtual tape.
- B. You are billed for the virtual tape data you store in Amazon Glacier and billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape.
- C. You are billed for the virtual tape data you store in Amazon S3 and billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape.
- D. You are billed for the virtual tape data you store in Amazon S3 and are billed for the size of the virtual tape.

Answer: B

Explanation: The AWS Storage Gateway is a service connecting an on-premises software appliance with cloud-based storage to provide seamless and secure integration between an organization's on-premises IT environment and AWS's storage infrastructure.

AWS Storage Gateway billing is as follows. Volume storage usage (per GB per month):

You are billed for the Cached volume data you store in Amazon S3. You are only billed for volume capacity you use, not for the size of the volume you create.

Snapshot Storage usage (per GB per month): You are billed for the snapshots your gateway stores in Amazon S3. These snapshots are stored and billed as Amazon EBS snapshots. Snapshots are incremental backups, reducing your storage charges. When taking a new snapshot, only the data that has changed since your last snapshot is stored.

Virtual Tape Library usage (per GB per month):

You are billed for the virtual tape data you store in Amazon S3. You are only billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape.

Virtual Tape Shelf usage (per GB per month):

You are billed for the virtual tape data you store in Amazon Glacier. You are only billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape.

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

NEW QUESTION 233

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

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

Answer: A

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

Reference: http://media.amazonwebservices.com/AWS_Building_Fault_Tolerant_Applications.pdf

NEW QUESTION 234

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 238

What is the default maximum number of Access Keys per user?

- A. 10
- B. 15
- C. 2
- D. 20

Answer: C

Explanation: The default maximum number of Access Keys per user is 2.

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

NEW QUESTION 241

What is the network performance offered by the c4.8xlarge instance in Amazon EC2?

- A. 20 Gigabit
- B. 10 Gigabit
- C. Very High but variable
- D. 5 Gigabit

Answer: B

Explanation: Networking performance offered by the c4.8xlarge instance is 10 Gigabit. Reference: <http://aws.amazon.com/ec2/instance-types/>

NEW QUESTION 242

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

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

Answer: C

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

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

The Resources section is the only section that is required.

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

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

NEW QUESTION 243

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

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

Answer: B

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

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

NEW QUESTION 245

How long does an AWS free usage tier EC2 last for?

- A. Forever
- B. 12 Months upon signup
- C. 1 Month upon signup
- D. 6 Months upon signup

Answer: B

Explanation: The AWS free usage tier will expire 12 months from the date you sign up. When your free usage expires or if your application use exceeds the free usage tiers, you simply pay the standard, pay-as-you-go service rates.

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

NEW QUESTION 250

Which of the following statements is true of tagging an Amazon EC2 resource?

- A. You don't need to specify the resource identifier while terminating a resource.
- B. You can terminate, stop, or delete a resource based solely on its tags.
- C. You can't terminate, stop, or delete a resource based solely on its tags.
- D. You don't need to specify the resource identifier while stopping a resource.

Answer: C

Explanation: You can assign tags only to resources that already exist. You can't terminate, stop, or delete a resource based solely on its tags; you must specify the resource identifier.

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

NEW QUESTION 251

You have been asked to tighten up the password policies in your organization after a serious security breach, so you need to consider every possible security measure. Which of the following is not an account password policy for IAM Users that can be set?

- A. Force IAM users to contact an account administrator when the user has allowed his or her password to expire.
- B. A minimum password length.
- C. Force IAM users to contact an account administrator when the user has entered his password incorrectly.
- D. Prevent IAM users from reusing previous password

Answer: C

Explanation: IAM users need passwords in order to access the AWS Management Console. (They do not need passwords if they will access AWS resources programmatically by using the CLI, AWS SDKs, or the APIs.)

You can use a password policy to do these things: Set a minimum password length.

Require specific character types, including uppercase letters, lowercase letters, numbers, and non-alphanumeric characters. Be sure to remind your users that passwords are case sensitive. Allow all IAM users to change their own passwords.

Require IAM users to change their password after a specified period of time (enable password expiration). Prevent IAM users from reusing previous passwords.

Force IAM users to contact an account administrator when the user has allowed his or her password to expire.

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

NEW QUESTION 254

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

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

Answer: A

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

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

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

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

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

NEW QUESTION 256

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

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

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

Answer: D

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

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

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

NEW QUESTION 261

A user has created a CloudFormation stack. The stack creates AWS services, such as EC2 instances, ELB, AutoScaling, and RDS. While creating the stack it created EC2, ELB and AutoScaling but failed to create RDS. What will CloudFormation do in this scenario?

- A. Rollback all the changes and terminate all the created services
- B. It will wait for the user's input about the error and correct the mistake after the input
- C. CloudFormation can never throw an error after launching a few services since it verifies all the steps before launching
- D. It will warn the user about the error and ask the user to manually create RDS

Answer: A

Explanation: AWS CloudFormation is an application management tool which provides application modeling, deployment, configuration, management and related activities. The AWS CloudFormation stack is a collection of AWS resources which are created and managed as a single unit when AWS CloudFormation instantiates a template. If any of the services fails to launch, CloudFormation will rollback all the changes and terminate or delete all the created services.

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

NEW QUESTION 263

In Amazon EC2, what is the limit of Reserved Instances per Availability Zone each month?

- A. 5
- B. 20
- C. 50
- D. 10

Answer: B

Explanation: There are 20 Reserved Instances per Availability Zone in each month.

Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html

NEW QUESTION 265

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

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

Answer: B

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

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

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

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

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

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

NEW QUESTION 269

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 274

After setting up some EC2 instances you now need to set up a monitoring solution to keep track of these instances and to send you an email when the CPU hits a certain threshold. Which statement below best describes what thresholds you can set to trigger a CloudWatch Alarm?

- A. Set a target value and choose whether the alarm will trigger when the value is greater than (>), greater than or equal to (>=), less than (<), or less than or equal to (<=) that value.
- B. Thresholds need to be set in IAM not CloudWatch
- C. Only default thresholds can be set you can't choose your own thresholds.
- D. Set a target value and choose whether the alarm will trigger when the value hits this threshold

Answer: A

Explanation: Amazon CloudWatch is a monitoring service for AWS cloud resources and the applications you run on AWS. You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, and set alarms.

When you create an alarm, you first choose the Amazon CloudWatch metric you want it to monitor. Next, you choose the evaluation period (e.g., five minutes or one hour) and a statistical value to measure (e.g., Average or Maximum).

To set a threshold, set a target value and choose whether the alarm will trigger when the value is greater than (>), greater than or equal to (>=), less than (<), or less than or equal to (<=) that value.

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

NEW QUESTION 277

After moving an E-Commerce website for a client from a dedicated server to AWS you have also set up auto scaling to perform health checks on the instances in your group and replace instances that fail these checks. Your client has come to you with his own health check system that he wants you to use as it has proved to be very useful prior to his site running on AWS. What do you think would be an appropriate response to this given all that you know about auto scaling?

- A. It is not possible to implement your own health check system
- B. You need to use AWS's health check system.
- C. It is not possible to implement your own health check system due to compatibility issues.
- D. It is possible to implement your own health check system and then send the instance's health information directly from your system to Cloud Watch.
- E. It is possible to implement your own health check system and then send the instance's health information directly from your system to Cloud Watch but only in the US East (
- F. Virginia) region.

Answer: C

Explanation: Auto Scaling periodically performs health checks on the instances in your group and replaces instances that fail these checks. By default, these health checks use the results of EC2 instance status checks to determine the health of an instance. If you use a load balancer with your Auto Scaling group, you can optionally choose to include the results of Elastic Load Balancing health checks.

Auto Scaling marks an instance unhealthy if the calls to the Amazon EC2 action DescribeInstanceStatus returns any other state other than running, the system status shows impaired, or the calls to Elastic Load Balancing action DescribeInstanceHealth returns OutOfService in the instance state field.

After an instance is marked unhealthy because of an Amazon EC2 or Elastic Load Balancing health check, it is scheduled for replacement.

You can customize the health check conducted by your Auto Scaling group by specifying additional checks or by having your own health check system and then sending the instance's health information directly from your system to Auto Scaling.

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

NEW QUESTION 281

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 286

You have just discovered that you can upload your objects to Amazon S3 using Multipart Upload API. You start to test it out but are unsure of the benefits that it would provide. Which of the following is not a benefit of using multipart uploads?

- A. You can begin an upload before you know the final object size.
- B. Quick recovery from any network issues.
- C. Pause and resume object uploads.
- D. It's more secure than normal upload.

Answer: D

Explanation: Multipart upload in Amazon S3 allows you to upload a single object as a set of parts. Each part is a contiguous portion of the object's data. You can upload these object parts independently and in any order.

If transmission of any part fails, you can re-transmit that part without affecting other parts. After all parts of your object are uploaded, Amazon S3 assembles these parts and creates the object. In general, when

your object size reaches 100 MB, you should consider using multipart uploads instead of uploading the object in a single operation.

Using multipart upload provides the following advantages:

Improved throughput—You can upload parts in parallel to improve throughput.

Quick recovery from any network issues—Smaller part size minimizes the impact of restarting a failed upload due to a network error.

Pause and resume object uploads—You can upload object parts over time. Once you initiate a multipart upload there is no expiry; you must explicitly complete or abort the multipart upload.

Begin an upload before you know the final object size—You can upload an object as you are creating it. Reference:

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

NEW QUESTION 288

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 292

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 297

Mike is appointed as Cloud Consultant in Netcrak Inc. Netcrak has the following VPCs set-up in the US East Region:

A VPC with CIDR block 10.10.0.0/16, a subnet in that VPC with CIDR block 10.10.1.0/24 A VPC with CIDR block 10.40.0.0/16, a subnet in that VPC with CIDR block 10.40.1.0/24

Netcrak Inc is trying to establish network connection between two subnets, a subnet with CIDR block 10.10.1.0/24 and another subnet with CIDR block 10.40.1.0/24. Which one of the following solutions should Mke recommend to Netcrak Inc?

- A. Create 2 Virtual Private Gateways and configure one with each VPC.
- B. Create one EC2 instance in each subnet, assign Elastic IPs to both instances, and configure a set up Site-to-Site VPN connection between both EC2 instances.
- C. Create a VPC Peering connection between both VPCs.
- D. Create 2 Internet Gateways, and attach one to each VP

Answer: C

Explanation: A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IP addresses. EC2 instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, or with a VPC in another AWS account within a single region.

AWS uses the existing infrastructure of a VPC to create a VPC peering connection; it is neither a gateway nor a VPN connection, and does not rely on a separate piece of physical hardware.

Reference: <http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-peering.htm>

NEW QUESTION 300

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

AWS-Solution-Architect-Associate Practice Exam Features:

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

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-Solution-Architect-Associate Practice Test Here](#)