

Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate

<https://www.2passeasy.com/dumps/AWS-Solution-Architect-Associate/>



NEW QUESTION 1

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 2

A user is launching an EC2 instance in the US East region. Which of the below mentioned options is recommended by AWS with respect to the selection of the availability zone?

- A. Always select the AZ while launching an instance
- B. Always select the US-East-1-a zone for HA
- C. Do not select the AZ; instead let AWS select the AZ
- D. The user can never select the availability zone while launching an instance

Answer: C

Explanation:

When launching an instance with EC2, AWS recommends not to select the availability zone (AZ). AWS specifies that the default Availability Zone should be accepted. This is because it enables AWS to select the best Availability Zone based on the system health and available capacity. If the user launches additional instances, only then an Availability Zone should be specified. This is to specify the same or different AZ from the running instances.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html>

NEW QUESTION 3

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 4

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 5

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 6

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 7

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 8

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 9

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 10

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 10

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 14

You've created your first load balancer and have registered your EC2 instances with the load balancer. Elastic Load Balancing routinely performs health checks on all the registered EC2 instances and automatically distributes all incoming requests to the DNS name of your load balancer across your registered, healthy EC2 instances. By default, the load balancer uses the _ protocol for checking the health of your instances.

- A. HTTPS
- B. HTTP
- C. ICMP
- D. IPv6

Answer: B

Explanation:

In Elastic Load Balancing a health configuration uses information such as protocol, ping port, ping path (URL), response timeout period, and health check interval to determine the health state of the instances registered with the load balancer.

Currently, HTTP on port 80 is the default health check. Reference:

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

NEW QUESTION 18

is a fast, filexible, fully managed push messaging service.

- A. Amazon SNS
- B. Amazon SES
- C. Amazon SQS
- D. Amazon FPS

Answer: A

Explanation:

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

Reference: http://aws.amazon.com/sns/?nc1=h_l2_as

NEW QUESTION 23

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 27

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 AMIs on Amazon EC2 and EBS

Answer: C

Explanation:

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

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

NEW QUESTION 30

In Amazon EC2, while sharing an Amazon EBS snapshot, can the snapshots with AWS IAM 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 35

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/cli/>

NEW QUESTION 36

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 37

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

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

Answer: C

Explanation:

Amazon Glacier supports various vault operations.

A vault inventory refers to the list of archives in a vault. For each archive in the list, the inventory provides archive information such as archive ID, creation date, and size. Amazon Glacier updates the vault inventory approximately once a day, starting on the day the first archive is uploaded to the vault. A vault inventory must exist for you to be able to download it.

Downloading a vault inventory is an asynchronous operation. You must first initiate a job to download the inventory. After receiving the job request, Amazon Glacier prepares your inventory for download. After the job completes, you can download the inventory data.

Given the asynchronous nature of the job, you can use Amazon Simple Notification Service (Amazon SNS) notifications to notify you when the job completes. You can specify an Amazon SNS topic for each individual job request or configure your vault to send a notification when specific vault events occur. Amazon Glacier prepares an inventory for each vault periodically, every 24 hours. If there have been no archive additions or deletions to the vault since the last inventory, the inventory date is not updated. When you initiate a job for a vault inventory, Amazon Glacier returns the last inventory it generated, which is a point-in-time snapshot and not real-time data. You might not find it useful to retrieve vault inventory for each archive upload. However, suppose you maintain a database on the client-side associating metadata about the archives you upload to Amazon Glacier. Then, you might find the vault inventory useful to reconcile information in your database with the actual vault inventory.

Reference: <http://docs.aws.amazon.com/amazonglacier/latest/dev/working-with-vaults.html>

NEW QUESTION 39

You are in the process of building an online gaming site for a client and one of the requirements is that it must be able to process vast amounts of data easily. Which AWS Service would be very helpful in processing all this data?

- A. Amazon S3
- B. AWS Data Pipeline
- C. AWS Direct Connect
- D. Amazon EMR

Answer: D

Explanation:

Managing and analyzing high data volumes produced by online games platforms can be difficult. The back-end infrastructures of online games can be challenging to maintain and operate. Peak usage periods, multiple players, and high volumes of write operations are some of the most common problems that operations teams face.

Amazon Elastic MapReduce (Amazon EMR) is a service that processes vast amounts of data easily. Input data can be retrieved from web server logs stored on Amazon S3 or from player data stored in Amazon DynamoDB tables to run analytics on player behavior, usage patterns, etc. Those results can be stored again on Amazon S3, or inserted in a relational database for further analysis with classic business intelligence tools.

Reference: http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_games_10.pdf

NEW QUESTION 42

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 47

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 48

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 50

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

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

Answer: A

Explanation:

The user can join multiple provisioned IOPS volumes together in a RAID 1 configuration to achieve better fault tolerance. RAID 1 does not provide a write performance improvement; it requires more bandwidth than non-RAID configurations since the data is written simultaneously to multiple volumes.

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

NEW QUESTION 55

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 57

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

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

Answer: A

Explanation:

After you launch an instance in EC2-Classic, you can't change its security groups. However, you can add rules to or remove rules from a security group, and those changes are automatically applied to all instances that are associated with the security group.

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

NEW QUESTION 58

A user has created photo editing software and hosted it on EC2. The software accepts requests from the user about the photo format and resolution and sends a message to S3 to enhance the picture accordingly. Which of the below mentioned AWS services will help make a scalable software with the AWS infrastructure in this scenario?

- A. AWS Simple Notification Service
- B. AWS Simple Queue Service
- C. AWS Elastic Transcoder
- D. AWS Glacier

Answer: B

Explanation:

Amazon Simple Queue Service (SQS) is a fast, reliable, scalable, and fully managed message queuing service. SQS provides a simple and cost-effective way to decouple the components of an application. The user can configure SQS, which will decouple the call between the EC2 application and S3. Thus, the application does not keep waiting for S3 to provide the data.

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

NEW QUESTION 61

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 64

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 67

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 68

A user comes to you and wants access to Amazon CloudWatch but only wants to monitor a specific LoadBalancer. Is it possible to give him access to a specific set of instances or a specific LoadBalancer?

- A. No because you can't use IAM to control access to CloudWatch data for specific resources.
- B. Yes
- C. You can use IAM to control access to CloudWatch data for specific resources.
- D. No because you need to be Sysadmin to access CloudWatch data.
- E. Yes
- F. Any user can see all CloudWatch data and needs no access right

Answer: A

Explanation:

Amazon CloudWatch integrates with AWS Identity and Access Management (IAM) so that you can specify which CloudWatch actions a user in your AWS Account can perform. For example, you could create an IAM policy that gives only certain users in your organization permission to use GetMetricStatistics. They could then use the action to retrieve data about your cloud resources.

You can't use IAM to control access to CloudWatch data for specific resources. For example, you can't give a user access to CloudWatch data for only a specific set of instances or a specific LoadBalancer. Permissions granted using IAM cover all the cloud resources you use with CloudWatch. In addition, you can't use IAM roles with the Amazon CloudWatch command line tools.

Using Amazon CloudWatch with IAM doesn't change how you use CloudWatch. There are no changes to CloudWatch actions, and no new CloudWatch actions related to users and access control.

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

NEW QUESTION 73

After setting up an EC2 security group with a cluster of 20 EC2 instances, you find an error in the security group settings. You quickly make changes to the security group settings. When will the changes to the settings be effective?

- A. The settings will be effective immediately for all the instances in the security group.
- B. The settings will be effective only when all the instances are restarted.
- C. The settings will be effective for all the instances only after 30 minutes.
- D. The settings will be effective only for the new instances added to the security group

Answer: A

Explanation:

Amazon Redshift applies changes to a cluster security group immediately. So if you have associated the cluster security group with a cluster, inbound cluster access rules in the updated cluster security group apply immediately.

Reference: <http://docs.aws.amazon.com/redshift/latest/mgmt/working-with-security-groups.html>

NEW QUESTION 76

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 79

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

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

Answer: B

Explanation:

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

infrastructure at very low costs.
AWS charges less where their costs are less.
For example, their costs are lower in the US Standard Region than in the US West (Northern California) Region.
Reference: <https://aws.amazon.com/s3/faqs/>

NEW QUESTION 81

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 86

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 89

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 93

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 95

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 96

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 97

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 102

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

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

Answer: A

Explanation:

Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. Amazon DynamoDB enables customers to offload the administrative burdens of operating and scaling distributed databases to AWS, so they don't have to worry about hardware provisioning, setup and configuration, replication, software patching, or cluster scaling.

Today's web-based applications generate and consume massive amounts of data. For example, an online game might start out with only a few thousand users and a light database workload consisting of 10 writes per second and 50 reads per second. However, if the game becomes successful, it may rapidly grow to millions of users and generate tens (or even hundreds) of thousands of writes and reads per second. It may also create terabytes or more of data per day. Developing your applications against Amazon DynamoDB enables you to start small and simply dial-up your request capacity for a table as your requirements scale, without incurring downtime. You pay highly cost-efficient rates for the request capacity you provision, and let Amazon DynamoDB do the work over partitioning your data and traffic over sufficient server capacity to meet your needs. Amazon DynamoDB does the database management and administration, and you simply store and request your data. Automatic replication and failover provides built-in fault tolerance, high availability, and data durability. Amazon DynamoDB gives you the peace of mind that your database is fully managed and can grow with your application requirements.

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

NEW QUESTION 107

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

- A. Are stateful: Return traffic is automatically allowed, regardless of any rules.
- B. Evaluate all rules before deciding whether to allow traffic.

- C. Support allow rules and deny rules.
- D. Operate at the instance level (first layer of defense).

Answer: C

Explanation:

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

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

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

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

NEW QUESTION 110

You are using Amazon SES as an email solution but are unsure of what its limitations are. Which statement below is correct in regards to that?

- A. New Amazon SES users who have received production access can send up to 1,000 emails per 24-hour period, at a maximum rate of 10 emails per second.
- B. Every Amazon SES sender has a the same set of sending limits
- C. Sending limits are based on messages rather than on recipients
- D. Every Amazon SES sender has a unique set of sending limits

Answer: D

Explanation:

Amazon Simple Email Service (Amazon SES) is a highly scalable and cost-effective email-sending service for businesses and developers. Amazon SES eliminates the complexity and expense of building an in-house email solution or licensing, installing, and operating a third-party email service for this type of email communication.

Every Amazon SES sender has a unique set of sending limits, which are calculated by Amazon SES on an ongoing basis:

Sending quota — the maximum number of emails you can send in a 24-hour period. Maximum send rate — the maximum number of emails you can send per second.

New Amazon SES users who have received production access can send up to 10,000 emails per 24-hour period, at a maximum rate of 5 emails per second.

Amazon SES automatically adjusts these limits upward, as long as you send high-quality email. If your existing quota is not adequate for your needs and the system has not automatically increased your quota, you can submit an SES Sending Quota Increase case at any time.

Sending limits are based on recipients rather than on messages. You can check your sending limits at any time by using the Amazon SES console.

Note that if your email is detected to be of poor or QUESTION able quality (e.g., high complaint rates, high bounce rates, spam, or abusive content), Amazon SES might temporarily or permanently reduce your permitted send volume, or take other action as AWS deems appropriate.

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

NEW QUESTION 115

Amazon Elastic Load Balancing is used to manage traffic on a fleet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits.

Which of the following is not an advantage of ELB over an on-premise load balancer?

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

Answer: A

Explanation:

Amazon Elastic Load Balancing is used to manage traffic on a fleet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits:

Takes over the encryption and decryption work from the Amazon EC2 instances and manages it centrally on the load balancer

Offers clients a single point of contact, and can also serve as the first line of defense against attacks on your network

When used in an Amazon VPC, supports creation and management of security groups associated with your Elastic Load Balancing to provide additional networking and security options

Supports end-to-end traffic encryption using TLS (previously SSL) on those networks that use secure HTTP (HTTPS) connections. When TLS is used, the TLS server certificate used to terminate client connections can be managed centrally on the load balancer, rather than on every individual instance. Reference:

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

NEW QUESTION 117

You have set up an S3 bucket with a number of images in it and you have decided that you want anybody to be able to access these images, even anonymous users. To accomplish this you create a bucket policy. You will need to use an Amazon S3 bucket policy that specifies a in the principal element, which means anyone can access the bucket.

- A. hash tag (#)
- B. anonymous user
- C. wildcard (*)
- D. S3 user

Answer: C

Explanation:

You can use the AWS Policy Generator to create a bucket policy for your Amazon S3 bucket. You can then use the generated document to set your bucket policy by using the Amazon S3 console, by a number of third-party tools, or via your application.

You use an Amazon S3 bucket policy that specifies a wildcard (*) in the principal element, which means anyone can access the bucket. With anonymous access, anyone (including users without an AWS account) will be able to access the bucket.

Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/iam-troubleshooting.htm#d0e20565>

NEW QUESTION 122

In Amazon CloudFront, if you use Amazon EC2 instances and other custom origins with CloudFront, it is recommended to .

- A. not use Elastic Load Balancing
- B. restrict Internet communication to private instances while allowing outgoing traffic
- C. enable access key rotation for CloudWatch metrics
- D. specify the URL of the load balancer for the domain name of your origin server

Answer: D

Explanation:

In Amazon CloudFront, you should use an Elastic Load Balancing load balancer to handle traffic across multiple Amazon EC2 instances and to isolate your application from changes to Amazon EC2 instances. When you create your CloudFront distribution, specify the URL of the load balancer for the domain name of your origin server.

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

NEW QUESTION 124

Are penetration tests allowed as long as they are limited to the customer's instances?

- A. Yes, they are allowed but only for selected regions.
- B. No, they are never allowed.
- C. Yes, they are allowed without any permission.
- D. Yes, they are allowed but only with approval.

Answer: D

Explanation:

Penetration tests are allowed after obtaining permission from AWS to perform them. Reference: <http://aws.amazon.com/security/penetration-testing/>

NEW QUESTION 125

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

- A. Luna Restore HSNI.
- B. Luna Backup HSM.
- C. Luna HSNI.
- D. Luna SA HSM.

Answer: B

Explanation:

In AWS CloudHSM, you can perform a remote backup/restore of a Luna SA partition if you have purchased a Luna Backup HSM.

Reference: <http://docs.aws.amazon.com/cloudhsm/latest/userguide/cloud-hsm-backup-restore.html>

NEW QUESTION 129

An EC2 instance is connected to an ENI (Elastic Network Interface) in one subnet. What happens when you attach an ENI of a different subnet to this EC2 instance?

- A. The EC2 instance follows the rules of the older subnet
- B. The EC2 instance follows the rules of both the subnets
- C. Not possible, cannot be connected to 2 ENIs
- D. The EC2 instance follows the rules of the newer subnet

Answer: B

Explanation:

AWS allows you create an elastic network interface (ENI), attach an ENI to an EC2 instance, detach an ENI from an EC2 instance and attach this ENI to another EC2 instance. The attributes of a network traffic follow the ENI which is attached to an EC2 instance or detached from an EC2 instance. When you move an ENI from one EC2 instance to another, network traffic is redirected to the new EC2 instance. You can create and attach additional ENIs to an EC2 instance.

Attaching multiple network interfaces (ENIs) to an EC2 instance is useful to: Create a management network.

Use network and security appliances in your VPC.

Create dual-homed instances with workloads/roles on distinct subnets Create a low-budget, high-availability solution.

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

NEW QUESTION 134

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

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

Answer: D

Explanation:

BitTorrent is an open, peer-to-peer protocol for distributing files. You can use the BitTorrent protocol to retrieve any publicly-accessible object in Amazon S3.

Amazon S3 supports the BitTorrent protocol so that developers can save costs when distributing content at high scale. Amazon S3 is useful for simple, reliable

storage of any data. The default distribution mechanism for Amazon S3 data is via client/server download. In client/server distribution, the entire object is transferred point-to-point from Amazon S3 to every authorized user who requests that object. While client/server delivery is appropriate for a wide variety of use cases, it is not optimal for everybody. Specifically, the costs of client/server distribution increase linearly as the number of users downloading objects increases. This can make it expensive to distribute popular objects.

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

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

NEW QUESTION 138

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 143

George has launched three EC2 instances inside the US-East-1a zone with his AWS account. Ray has launched two EC2 instances in the US-East-1a zone with his AWS account. Which of the below mentioned statements will help George and Ray understand the availability zone (AZ) concept better?

- A. All the instances of George and Ray can communicate over a private IP with a minimal cost
- B. The US-East-1a region of George and Ray can be different availability zones
- C. All the instances of George and Ray can communicate over a private IP without any cost
- D. The instances of George and Ray will be running in the same data centre

Answer: B

Explanation:

Each AWS region has multiple, isolated locations known as Availability Zones. To ensure that the AWS resources are distributed across the Availability Zones for a region, AWS independently maps the Availability Zones to identifiers for each account. In this case the Availability Zone US-East-1a where George's EC2 instances are running might not be the same location as the US-East-1a zone of Ray's EC2 instances. There is no way for the user to coordinate the Availability Zones between accounts.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html>

NEW QUESTION 144

A user has created an ELB with Auto Scaling. Which of the below mentioned offerings from ELB helps the user to stop sending new requests traffic from the load balancer to the EC2 instance when the instance is being deregistered while continuing in-flight requests?

- A. ELB sticky session
- B. ELB deregistration check
- C. ELB auto registration Off
- D. ELB connection draining

Answer: D

Explanation:

The Elastic Load Balancer connection draining feature causes the load balancer to stop sending new requests to the back-end instances when the instances are deregistering or become unhealthy, while ensuring that in-flight requests continue to be served.

Reference:

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

NEW QUESTION 147

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 148

You have been asked to set up monitoring of your network and you have decided that Cloudwatch would be the best service to use. Amazon CloudWatch monitors your Amazon Web Services (AWS) resources and the applications you run on AWS in real-time. You can use CloudWatch to collect and track metrics, which are the variables you want to measure for your resources and applications. Which of the following items listed can AWS Cloudwatch monitor?

- A. Log files your applications generate.
- B. All of the items listed on this page.
- C. System-wide visibility into resource utilization, application performance, and operational health.
- D. Custom metrics generated by your applications and services .

Answer: B

Explanation:

Amazon CloudWatch can monitor AWS resources such as Amazon EC2 instances, Amazon DynamoDB tables, and Amazon RDS DB instances, as well as custom metrics generated by your applications and services, and any log files your applications generate. You can use Amazon CloudWatch to gain system-wide visibility into resource utilization, application performance, and operational health. You can use these insights to react and keep your application running smoothly.

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

NEW QUESTION 149

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

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

Answer: A

Explanation:

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

Reference:

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

NEW QUESTION 151

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 152

Which of the following statements is NOT true about using Elastic IP Address (EIP) in EC2-Classic and EC2-VPC platforms?

- A. In the EC2-VPC platform, the Elastic IP Address (EIP) does not remain associated with the instance when you stop it.
- B. In the EC2-Classic platform, stopping the instance disassociates the Elastic IP Address (EIP) from it.
- C. In the EC2-VPC platform, if you have attached a second network interface to an instance, when you disassociate the Elastic IP Address (EIP) from that instance, a new public IP address is not assigned to the instance automatically; you'll have to associate an EIP with it manually.
- D. In the EC2-Classic platform, if you disassociate an Elastic IP Address (EIP) from the instance, the instance is automatically assigned a new public IP address within a few minutes.

Answer: A

Explanation:

In the EC2-Classic platform, when you associate an Elastic IP Address (EIP) with an instance, the instance's current public IP address is released to the EC2-Classic public IP address pool. If you disassociate an EIP from the instance, the instance is automatically assigned a new public IP address within a few minutes. In addition, stopping the instance also disassociates the EIP from it.

But in the EC2-VPC platform, when you associate an EIP with an instance in a default Virtual Private Cloud (VPC), or an instance in which you assigned a public IP to the eth0 network interface during launch, its current public IP address is released to the EC2-VPC public IP address pool. If you disassociate an EIP from the instance, the instance is automatically assigned a new public IP address within a few minutes. However, if you have attached a second network interface to the instance, the instance is not automatically assigned a new public IP address; you'll have to associate an EIP with it manually. The EIP remains associated with the instance when you stop it.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html>

NEW QUESTION 153

A user has launched one EC2 instance in the US West region. The user wants to access the RDS instance launched in the US East region from that EC2 instance. How can the user configure the access for that EC2 instance?

- A. Configure the IP range of the US West region instance as the ingress security rule of RDS
- B. It is not possible to access RDS of the US East region from the US West region
- C. Open the security group of the US West region in the RDS security group's ingress rule
- D. Create an IAM role which has access to RDS and launch an instance in the US West region with it

Answer: A

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.

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

NEW QUESTION 157

Company B is launching a new game app for mobile devices. Users will log into the game using their existing social media account to streamline data capture. Company B would like to directly save player data and scoring information from the mobile app to a DynamoDB table named Score Data. When a user saves their game the progress data will be stored to the Game state 53 bucket. What is the best approach for storing data to DynamoDB and S3?

- A. Use an EC2 Instance that is launched with an EC2 role providing access to the Score Data DynamoDB table and the Game State S3 bucket that communicates with the mobile app via web services.
- B. Use temporary security credentials that assume a role providing access to the Score Data DynamoDB table and the Game State S3 bucket using web identity federation.
- C. Use Login with Amazon allowing users to sign in with an Amazon account providing the mobile app with access to the Score Data DynamoDB table and the Game State S3 bucket.
- D. Use an IAM user with access credentials assigned a role providing access to the Score Data DynamoDB table and the Game State S3 bucket for distribution with the mobile app.

Answer: B

Explanation:

Web Identity Federation

Imagine that you are creating a mobile app that accesses AWS resources, such as a game that runs on a mobile device and stores player and score information using Amazon S3 and DynamoDB. When you write such an app, you'll make requests to AWS services that must be signed with an AWS access key. However, we strongly recommend that you do not embed or distribute long-term AWS credentials with apps that a user downloads to a device, even in an encrypted store. Instead, build your app so that it requests temporary AWS security credentials dynamically when needed using web identity federation. The supplied temporary credentials map to an AWS role that has only the permissions needed to perform the tasks required by the mobile app.

With web identity federation, you don't need to create custom sign-in code or manage your own user identities. Instead, users of your app can sign in using a well-known identity provider (IdP) - such as Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible IdP, receive an authentication token, and then exchange that token for temporary security credentials in AWS that map to an IAM role with permissions to use the resources in your AWS account. Using an IdP helps you keep your AWS account secure, because you don't have to embed and distribute long-term security credentials with your application.

For most scenarios, we recommend that you use Amazon Cognito because it acts as an identity broker and does much of the federation work for you. For details, see the following section, [Using Amazon Cognito for Mobile Apps](#).

If you don't use Amazon Cognito, then you must write code that interacts with a web IdP (Login with Amazon, Facebook, Google, or any other OIDC-compatible IdP) and then calls the Assume Role With Web Identity API to trade the authentication token you get from those IdPs for AWS temporary security credentials. If you have already used this approach for existing apps, you can continue to use it.

Using Amazon Cognito for Mobile Apps

The preferred way to use web identity federation is to use Amazon Cognito. For example, Adele the developer is building a game for a mobile device where user data such as scores and profiles is stored in Amazon S3 and Amazon DynamoDB. Adele could also store this data locally on the device and use Amazon Cognito to keep it synchronized across devices. She knows that for security and maintenance reasons, long-term AWS security credentials should not be distributed with the game. She also knows that the game might have a large number of users. For all of these reasons, she does not want to create new user identities in IAM for each player. Instead, she builds the game so that users can sign in using an identity that they've already established with a well-known identity provider, such as Login with Amazon, Facebook, Google, or any OpenID Connect (OIDC)-compatible identity provider.

Her game can take advantage of the authentication mechanism from one of these providers to validate the user's identity.

To enable the mobile app to access her AWS resources, Adele first registers for a developer ID with her chosen IdPs. She also configures the application with each of these providers. In her AWS account that contains the Amazon S3 bucket and DynamoDB table for the game, Adele uses Amazon Cognito to create IAM roles that precisely define permissions that the game needs. If she is using an OIDC IdP, she also creates an IAM OIDC identity provider entity to establish trust between her AWS account and the IdP.

In the app's code, Adele calls the sign-in interface for the IdP that she configured previously. The IdP handles all the details of letting the user sign in, and the app gets an OAuth access token or OIDC ID token from the provider. Adele's app can trade this authentication information for a set of temporary security credentials that consist of an AWS access key ID, a secret access key, and a session token.

The app can then use these credentials to access web services offered by AWS. The app is limited to the permissions that are defined in the role that it assumes. The following figure shows a simplified flow for how this might work, using Login with Amazon as the IdP.

For Step 2, the app can also use Facebook, Google, or any OIDC-compatible identity provider, but that's not shown here.

Sample workflow using Amazon Cognito to federate users for a mobile application

A customer starts your app on a mobile device. The app asks the user to sign in. The app uses Login with Amazon resources to accept the user's credentials. The app uses Cognito APIs to exchange the Login with Amazon ID token for a Cognito token. The app requests temporary security credentials from AWS STS, passing the Cognito token.

The temporary security credentials can be used by the app to access any AWS resources required by the app to operate. The role associated with the temporary security credentials and its assigned policies determines what can be accessed.

Use the following process to configure your app to use Amazon Cognito to authenticate users and give your app access to AWS resources. For specific steps to accomplish this scenario, consult the documentation for Amazon Cognito.

(Optional) Sign up as a developer with Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible identity provider and configure one or more apps with the provider. This step is optional because Amazon Cognito also supports unauthenticated (guest) access for your users.

Go to Amazon Cognito in the AWS IAM Management Console. Use the Amazon Cognito wizard to create an identity pool, which is a container that Amazon Cognito uses to keep end user identities organized for your apps. You can share identity pools between apps. When you set up an identity pool, Amazon Cognito creates one or two IAM roles (one for authenticated identities, and one for unauthenticated "guest" identities) that define permissions for Amazon Cognito users.

Download and integrate the AWS SDK for iOS or the AWS SDK for Android with your app, and import the files required to use Amazon Cognito.

Create an instance of the Amazon Cognito credentials provider, passing the identity pool ID, your AWS account number, and the Amazon Resource Name (ARN) of the roles that you associated with the identity pool. The Amazon Cognito wizard in the AWS Management Console provides sample code to help you get started.

When your app accesses an AWS resource, pass the credentials provider instance to the client object, which passes temporary security credentials to the client. The permissions for the credentials are based on the role or roles that you defined earlier.

NEW QUESTION 160

Your company plans to host a large donation website on Amazon Web Services (AWS). You anticipate a large and undetermined amount of traffic that will create many database writes. To be certain that you do not drop any writes to a database hosted on AWS. Which service should you use?

- A. Amazon RDS with provisioned IOPS up to the anticipated peak write throughput.
- B. Amazon Simple Queue Service (SQS) for capturing the writes and draining the queue to write to the database.
- C. Amazon ElastiCache to store the writes until the writes are committed to the database.
- D. Amazon DynamoDB with provisioned write throughput up to the anticipated peak write throughput

Answer: B

Explanation:

Amazon Simple Queue Service (Amazon SQS) offers a reliable, highly scalable hosted queue for storing messages as they travel between computers. By using Amazon SQS, developers can simply move data between distributed application components performing different tasks, without losing messages or requiring each component to be always available. Amazon SQS makes it easy to build a distributed, decoupled application, working in close conjunction with the Amazon Elastic Compute Cloud (Amazon EC2) and the other AWS infrastructure web services.

What can I do with Amazon SQS?

Amazon SQS is a web service that gives you access to a message queue that can be used to store messages while waiting for a computer to process them. This allows you to quickly build message queuing applications that can be run on any computer on the internet. Since Amazon SQS is highly scalable and you only pay for what you use, you can start small and grow your application as you wish, with no compromise on performance or reliability. This lets you focus on building sophisticated message-based applications, without worrying about how the messages are stored and managed.

You can use Amazon SQS with software applications in various ways. For example, you can: Integrate Amazon SQS with other AWS infrastructure web services to make applications more reliable and flexible.

Use Amazon SQS to create a queue of work where each message is a task that needs to be completed by a process. One or many computers can read tasks from the queue and perform them. Build a microservices architecture, using queues to connect your microservices.

Keep notifications of significant events in a business process in an Amazon SQS queue. Each event can have a corresponding message in a queue, and applications that need to be aware of the event can read and process the messages.

NEW QUESTION 161

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

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

Answer: E

NEW QUESTION 163

You have recently joined a startup company building sensors to measure street noise and air quality in urban areas. The company has been running a pilot deployment of around 100 sensors for 3 months each sensor uploads 1KB of sensor data every minute to a backend hosted on AWS.

During the pilot, you measured a peak of 10 IOPS on the database, and you stored an average of 3GB of sensor data per month in the database.

The current deployment consists of a load-balanced auto scaled ingestion layer using EC2 instances and a PostgreSQL RDS database with 500GB standard storage.

The pilot is considered a success and your CEO has managed to get the attention of some potential investors. The business plan requires a deployment of at least 1,000 sensors which needs to be supported by the backend. You also need to store sensor data for at least two years to be able to compare year over year improvements.

To secure funding, you have to make sure that the platform meets these requirements and leaves room for further scaling. Which setup will meet the requirements?

- A. Add an SQS queue to the ingestion layer to buffer writes to the RDS instance
- B. Ingest data into a DynamoDB table and move old data to a Redshift cluster
- C. Replace the RDS instance with a 6 node Redshift cluster with 96TB of storage
- D. Keep the current architecture but upgrade RDS storage to 3TB and IOPS provisioned IOPS

Answer: C

NEW QUESTION 167

You have been asked to design the storage layer for an application. The application requires disk

performance of at least 100,000 IOPS in addition, the storage layer must be able to survive the loss of an individual disk. EC2 instance, or Availability Zone without any data loss. The volume you provide must have a capacity of at least 3 TB. Which of the following designs will meet these objectives'?

- A. Instantiate a c3.8xlarge instance in us-east-1. Provision 4x1TB EBS volumes, attach them to the instance, and configure them as a single RAID 5 volume
- B. Ensure that EBS snapshots are performed every 15 minutes.
- C. Instantiate a c3.8xlarge instance in us-east-1. Provision 3x1TB EBS volumes, attach them to the Instance, and configure them as a single RAID 0 volume
- D. Ensure that EBS snapshots are performed every 15 minutes.
- E. Instantiate an i2.8xlarge instance in us-east-1

- F. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instance
- G. Provision 3x1TB EBS volumes, attach them to the instance, and configure them as a second RAID 0 volume
- H. Configure synchronous, block-level replication from the ephemeral-backed volume to the EBS-backed volume.
- I. Instantiate a c3.8xlarge instance in us-east-1. Provision an AWS Storage Gateway and configure it for 3 TB of storage and 100,000 IOP
- J. Attach the volume to the instance
- K. Instantiate an i2.8xlarge instance in us-east-1
- L. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instance
- M. Configure synchronous, block-level replication to an identically configured instance in us-east-1

Answer: C

NEW QUESTION 170

An ERP application is deployed across multiple AZs in a single region. In the event of failure, the Recovery Time Objective (RTO) must be less than 3 hours, and the Recovery Point Objective (RPO) must be 15 minutes. The customer realizes that data corruption occurred roughly 1.5 hours ago. What DR strategy could be used to achieve this RTO and RPO in the event of this kind of failure?

- A. Take hourly DB backups to S3, with transaction logs stored in S3 every 5 minutes.
- B. Use synchronous database master-slave replication between two availability zones.
- C. Take hourly DB backups to EC2 Instance store volumes with transaction logs stored in S3 every 5 minutes.
- D. Take 15 minute DB backups stored in Glacier with transaction logs stored in S3 every 5 minutes

Answer: A

NEW QUESTION 172

A customer has established an AWS Direct Connect connection to AWS. The link is up and routes are being advertised from the customer's end, however the customer is unable to connect from EC2 instances inside its VPC to servers residing in its datacenter. Which of the following options provide a viable solution to remedy this situation? (Choose 2 answers)

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

Answer: AC

NEW QUESTION 173

You have deployed a three-tier web application in a VPC with a CIDR block of 10.0.0.0/28. You initially deploy two web servers, two application servers, two database servers and one NAT instance for a total of seven EC2 instances. The web, application and database servers are deployed across two availability zones (AZs). You also deploy an ELB in front of the two web servers, and use Route53 for DNS. Web traffic gradually increases in the first few days following the deployment, so you attempt to double the number of instances in each tier of the application to handle the new load. Unfortunately, some of these new instances fail to launch.

Which of the following could be the root cause? (Choose 2 answers)

- A. AWS reserves the first and the last private IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances
- B. The Internet Gateway (IGW) of your VPC has scaled-up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches
- C. The ELB has scaled-up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches
- D. AWS reserves one IP address in each subnet's CIDR block for Route53 so you do not have enough addresses left to launch all of the new EC2 instances
- E. AWS reserves the first four and the last IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances

Answer: CE

NEW QUESTION 176

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

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

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

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

rtb-218bc449 rtb-238bc44b Associations:

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

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

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

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

Answer: A

NEW QUESTION 179

You are migrating a legacy client-server application to AWS. The application responds to a specific DNS domain (e.g. www.example.com) and has a 2-tier

architecture, with multiple application servers and a database server. Remote clients use TCP to connect to the application servers. The application servers need to know the IP address of the clients in order to function properly and are currently taking that information from the TCP socket. A Multi-AZ RDS MySQL instance will be used for the database. During the migration you can change the application code, but you have to file a change request. How would you implement the architecture on AWS in order to maximize scalability and high availability?

- A. File a change request to implement Alias Resource support in the application
- B. Use Route 53 Alias Resource Record to distribute load on two application servers in different AZs.
- C. File a change request to implement Latency Based Routing support in the application
- D. Use Route 53 with Latency Based Routing enabled to distribute load on two application servers in different AZs.
- E. File a change request to implement Cross-Zone support in the application
- F. Use an ELB with a TCP Listener and Cross-Zone Load Balancing enabled, two application servers in different AZs.
- G. File a change request to implement Proxy Protocol support in the application
- H. Use an ELB with a TCP Listener and Proxy Protocol enabled to distribute load on two application servers in different AZs.

Answer: D

NEW QUESTION 184

Your company has an on-premises multi-tier PHP web application, which recently experienced downtime due to a large burst in web traffic due to a company announcement. Over the coming days, you are expecting similar announcements to drive similar unpredictable bursts, and are looking to find ways to quickly improve your infrastructure's ability to handle unexpected increases in traffic.

The application currently consists of 2 tiers: a web tier which consists of a load balancer and several Linux Apache web servers as well as a database tier which hosts a Linux server hosting a MySQL database. Which scenario below will provide full site functionality, while helping to improve the ability of your application in the short timeframe required?

- A. Failover environment: Create an S3 bucket and configure it for website hosting
- B. Migrate your DNS to Route53 using zone file import, and leverage Route53 DNS failover to failover to the S3 hosted website.
- C. Hybrid environment: Create an AMI, which can be used to launch web servers in EC2. Create an Auto Scaling group, which uses the AMI to scale the web tier based on incoming traffic
- D. Leverage Elastic Load Balancing to balance traffic between on-premises web servers and those hosted in AWS.
- E. Offload traffic from on-premises environment: Setup a CloudFront distribution, and configure CloudFront to cache objects from a custom origin
- F. Choose to customize your object cache behavior, and select a TTL that objects should exist in cache.
- G. Migrate to AWS: Use VM Import/Export to quickly convert an on-premises web server to an AMI
- H. Create an Auto Scaling group, which uses the imported AMI to scale the web tier based on incoming traffic
- I. Create an RDS read replica and setup replication between the RDS instance and on-premises MySQL server to migrate the database.

Answer: C

NEW QUESTION 186

Your company produces customer commissioned one-of-a-kind skiing helmets combining high fashion with custom technical enhancements. Customers can show off their individuality on the ski slopes and have access to head-up-displays, GPS rear-view cams and any other technical innovation they wish to embed in the helmet.

The current manufacturing process is data rich and complex including assessments to ensure that the custom electronics and materials used to assemble the helmets are to the highest standards. Assessments are a mixture of human and automated assessments; you need to add a new set of assessment to model the failure modes of the custom electronics using GPUs with CUDA, across a cluster of servers with low latency networking.

What architecture would allow you to automate the existing process using a hybrid approach and ensure that the architecture can support the evolution of processes over time?

- A. Use AWS Data Pipeline to manage movement of data & meta-data and assessments. Use an auto-scaling group of G2 instances in a placement group.
- B. Use Amazon Simple Workflow (SWF) to manage assessments, movement of data & meta-data. Use an auto-scaling group of G2 instances in a placement group.
- C. Use Amazon Simple Workflow (SWF) to manage assessments, movement of data & meta-data. Use an auto-scaling group of C3 instances with SR-IOV (Single Root I/O Virtualization).
- D. Use AWS Data Pipeline to manage movement of data & meta-data and assessments. Use an auto-scaling group of C3 instances with SR-IOV (Single Root I/O virtualization).

Answer: B

NEW QUESTION 187

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

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

Answer: A

NEW QUESTION 189

An enterprise wants to use a third-party SaaS application. The SaaS application needs to have access to issue several API commands to discover Amazon EC2 resources running within the enterprise's account. The enterprise has internal security policies that require any outside access to their environment must conform to the principles of least privilege and there must be controls in place to ensure that the credentials used by the SaaS vendor cannot be used by any other third party. Which of the following would meet all of these conditions?

- A. From the AWS Management Console, navigate to the Security Credentials page and retrieve the access and secret key for your account.

- B. Create an IAM user within the enterprise account assign a user policy to the IAM user that allows only the actions required by the SaaS application create a new access and secret key for the user and provide these credentials to the 5aa5 provider.
- C. Create an IAM role for cross-account access allows the SaaS provider's account to assume the role and assign it a policy that allows only the actions required by the SaaS application.
- D. Create an IAM role for EC2 instances, assign it a policy that allows only the actions required for the SaaS application to work, provide the role ARM to the SaaS provider to use when launching their application instances.

Answer: C

Explanation:

Granting Cross-account Permission to objects It Does Not Own

In this example scenario, you own a bucket and you have enabled other AWS accounts to upload objects. That is, your bucket can have objects that other AWS accounts own.

Now, suppose as a bucket owner, you need to grant cross-account permission on objects, regardless of who the owner is, to a user in another account. For example, that user could be a billing application that needs to access object metadata. There are two core issues:

The bucket owner has no permissions on those objects created by other AWS accounts. So for the bucket owner to grant permissions on objects it does not own, the object owner, the AWS account that created the objects, must first grant permission to the bucket owner. The bucket owner can then delegate those permissions.

Bucket owner account can delegate permissions to users in its own account but it cannot delegate permissions to other AWS accounts, because cross-account delegation is not supported.

In this scenario, the bucket owner can create an AWS Identity and Access Management (IAM) role with permission to access objects, and grant another AWS account permission to assume the role temporarily enabling it to access objects in the bucket.

Background: Cross-Account Permissions and Using IAM Roles

IAM roles enable several scenarios to delegate access to your resources, and cross-account access is

one of the key scenarios. In this example, the bucket owner, Account A, uses an IAM role to temporarily delegate object access cross-account to users in another AWS account, Account C. Each IAM role you create has two policies attached to it:

A trust policy identifying another AWS account that can assume the role.

An access policy defining what permissions—for example, s3:GetObject—are allowed when someone assumes the role. For a list of permissions you can specify in a policy, see [Specifying Permissions in a Policy](#).

The AWS account identified in the trust policy then grants its user permission to assume the role. The user can then do the following to access objects:

Assume the role and, in response, get temporary security credentials. Using the temporary security credentials, access the objects in the bucket.

For more information about IAM roles, go to [Roles \(Delegation and Federation\)](#) in [IAM User Guide](#). The following is a summary of the walkthrough steps:

Account A administrator user attaches a bucket policy granting Account B conditional permission to upload objects.

Account A administrator creates an IAM role, establishing trust with Account C, so users in that account can access Account A. The access policy attached to the role limits what user in Account C can do when the user accesses Account A.

Account B administrator uploads an object to the bucket owned by Account A, granting full —control permission to the bucket owner.

Account C administrator creates a user and attaches a user policy that allows the user to assume the role. User in Account C first assumes the role, which returns the user temporary security credentials.

Using those temporary credentials, the user then accesses objects in the bucket.

For this example, you need three accounts. The following table shows how we refer to these accounts and the administrator users in these accounts. Per IAM guidelines (see [About Using an Administrator User to Create Resources and Grant Permissions](#)) we do not use the account root credentials in this walkthrough. Instead, you create an administrator user in each account and use those credentials in creating resources and granting them permissions

NEW QUESTION 190

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

How should the application use AWS credentials to access the S3 bucket securely?

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

Answer: C

NEW QUESTION 193

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

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

Answer: D

Explanation:

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

NEW QUESTION 195

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

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

Answer: ACD

NEW QUESTION 198

You require the ability to analyze a customer's clickstream data on a website so they can do behavioral analysis. Your customer needs to know what sequence of pages and ads their customer clicked on. This data will be used in real time to modify the page layouts as customers click through the site to increase stickiness and advertising click-through. Which option meets the requirements for capturing and analyzing this data?

- A. Log clicks in weblogs by URL store to Amazon S3, and then analyze with Elastic MapReduce
- B. Push web clicks by session to Amazon Kinesis and analyze behavior using Kinesis workers
- C. Write click events directly to Amazon Redshift and then analyze with SQL
- D. Publish web clicks by session to an Amazon SQS queue periodically drain these events to Amazon RDS and analyze with SQL

Answer: B

Explanation:

Reference: <http://www.slideshare.net/AmazonWebServices/aws-webcast-introduction-to-amazon-kinesis>

NEW QUESTION 202

An AWS customer runs a public blogging website. The site users upload two million blog entries a month. The average blog entry size is 200 KB. The access rate to blog entries drops to negligible 6 months after publication and users rarely access a blog entry 1 year after publication. Additionally, blog entries have a high update rate during the first 3 months following publication, this drops to no updates after 6 months. The customer wants to use CloudFront to improve his user's load times.

Which of the following recommendations would you make to the customer?

- A. Duplicate entries into two different buckets and create two separate CloudFront distributions where S3 access is restricted only to Cloud Front identity
- B. Create a CloudFront distribution with "US" Europe price class for US/ Europe users and a different CloudFront distribution with All Edge Locations' for the remaining users.
- C. Create a CloudFront distribution with S3 access restricted only to the CloudFront identity and partition the blog entry's location in S3 according to the month it was uploaded to be used with CloudFront behaviors.
- D. Create a CloudFront distribution with Restrict Viewer Access Forward Query string set to true and minimum TTL of 0.

Answer: C

NEW QUESTION 204

Your company is getting ready to do a major public announcement of a social media site on AWS. The website is running on EC2 instances deployed across multiple Availability Zones with a Multi-AZ RDS MySQL Extra Large DB Instance. The site performs a high number of small reads and writes per second and relies on an eventual consistency model. After comprehensive tests you discover that there is read contention on RDS MySQL. Which are the best approaches to meet these requirements? (Choose 2 answers)

- A. Deploy ElasticCache in-memory cache running in each availability zone
- B. Implement sharding to distribute load to multiple RDS MySQL instances
- C. Increase the RDS MySQL Instance size and Implement provisioned IOPS
- D. Add an RDS MySQL read replica in each availability zone

Answer: AC

NEW QUESTION 207

You are developing a new mobile application and are considering storing user preferences in AWS. This would provide a more uniform cross-device experience to users using multiple mobile devices to access the application. The preference data for each user is estimated to be 50KB in size. Additionally 5 million customers are expected to use the application on a regular basis. The solution needs to be cost-effective, highly available, scalable and secure, how would you design a solution to meet the above requirements?

- A. Setup an RDS MySQL instance in 2 availability zones to store the user preference data
- B. Deploy a public facing application on a server in front of the database to manage security and access credentials
- C. Setup a DynamoDB table with an item for each user having the necessary attributes to hold the user preference
- D. The mobile application will query the user preferences directly from the DynamoDB table
- E. Utilize STS
- F. Web Identity Federation, and DynamoDB Fine Grained Access Control to authenticate and authorize access.
- G. Setup an RDS MySQL instance with multiple read replicas in 2 availability zones to store the user preference data. The mobile application will query the user preferences from the read replica
- H. Leverage the MySQL user management and access privilege system to manage security and access credentials.
- I. Store the user preference data in S3. Setup a DynamoDB table with an item for each user and an item attribute pointing to the user's S3 object
- J. The mobile application will retrieve the S3 URL from DynamoDB and then access the S3 object directly utilize STS, Web identity Federation, and S3 ACLs to authenticate and authorize access.

Answer: B

NEW QUESTION 212

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

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

Answer: A

NEW QUESTION 213

You are looking to migrate your Development (Dev) and Test environments to AWS. You have decided to use separate AWS accounts to host each environment. You plan to link each account's bill to a Master AWS account using Consolidated Billing. To make sure you keep within budget you would like to implement a way for administrators in the Master account to have access to stop, delete and/or terminate resources in both the Dev and Test accounts. Identify which option will allow you to achieve this goal.

- A. Create IAM users in the Master account with full Admin permission.
- B. Create cross-account roles in the Dev and Test accounts that grant the Master account access to the resources in the account by inheriting permissions from the Master account.
- C. Create IAM users and a cross-account role in the Master account that grants full Admin permissions to the Dev and Test accounts.
- D. Create IAM users in the Master account. Create cross-account roles in the Dev and Test accounts that have full Admin permissions and grant the Master account access.
- E. Link the accounts using Consolidated Billing.
- F. This will give IAM users in the Master account access to resources in the Dev and Test accounts.

Answer: C

Explanation:

Bucket Owner Granting Cross-account Permission to objects It Does Not Own

In this example scenario, you own a bucket and you have enabled other AWS accounts to upload objects. That is, your bucket can have objects that other AWS accounts own.

Now, suppose as a bucket owner, you need to grant cross-account permission on objects, regardless of who the owner is, to a user in another account. For example, that user could be a billing application that needs to access object metadata. There are two core issues:

The bucket owner has no permissions on those objects created by other AWS accounts. So for the bucket owner to grant permissions on objects it does not own, the object owner, the AWS account that created the objects, must first grant permission to the bucket owner. The bucket owner can then delegate those permissions.

Bucket owner account can delegate permissions to users in its own account but it cannot delegate permissions to other AWS accounts, because cross-account delegation is not supported.

In this scenario, the bucket owner can create an AWS Identity and Access Management (IAM) role with permission to access objects, and grant another AWS account permission to assume the role temporarily enabling it to access objects in the bucket.

Background: Cross-Account Permissions and Using IAM Roles

IAM roles enable several scenarios to delegate access to your resources, and cross-account access is

one of the key scenarios. In this example, the bucket owner, Account A, uses an IAM role to temporarily delegate object access cross-account to users in another AWS account, Account C. Each IAM role you create has two policies attached to it:

A trust policy identifying another AWS account that can assume the role.

An access policy defining what permissions—for example, `s3:GetObject`—are allowed when someone assumes the role. For a list of permissions you can specify in a policy, see [Specifying Permissions in a Policy](#).

The AWS account identified in the trust policy then grants its user permission to assume the role. The user can then do the following to access objects:

Assume the role and, in response, get temporary security credentials. Using the temporary security credentials, access the objects in the bucket.

For more information about IAM roles, go to [Roles \(Delegation and Federation\)](#) in [IAM User Guide](#). The following is a summary of the walkthrough steps:

Account A administrator user attaches a bucket policy granting Account B conditional permission to upload objects.

Account A administrator creates an IAM role, establishing trust with Account C, so users in that account can access Account A. The access policy attached to the role limits what user in Account C can do when the user accesses Account A.

Account B administrator uploads an object to the bucket owned by Account A, granting full-control permission to the bucket owner.

Account C administrator creates a user and attaches a user policy that allows the user to assume the role.

User in Account C first assumes the role, which returns the user temporary security credentials. Using those temporary credentials, the user then accesses objects in the bucket.

For this example, you need three accounts. The following table shows how we refer to these accounts and the administrator users in these accounts. Per IAM guidelines (see [About Using an](#)

[Administrator User to Create Resources and Grant Permissions](#)) we do not use the account root

credentials in this walkthrough. Instead, you create an administrator user in each account and use those credentials in creating resources and granting them permissions.

NEW QUESTION 217

All Amazon EC2 instances are assigned two IP addresses at launch, out of which one can only be reached from within the Amazon EC2 network?

- A. Multiple IP address
- B. Public IP address
- C. Private IP address
- D. Elastic IP Address

Answer: C

NEW QUESTION 221

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

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

Answer: C

NEW QUESTION 223

If I write the below command, what does it do? `ec2-run ami-e3a5408a -n 20 -g appserver`

- A. Start twenty instances as members of appserver group.
- B. Creates 20 rules in the security group named appserver
- C. Terminate twenty instances as members of appserver group.
- D. Start 20 security groups

Answer: A

NEW QUESTION 228

Can I control if and when MySQL based RDS Instance is upgraded to new supported versions?

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

Answer: C

NEW QUESTION 229

What does Amazon Elastic Beanstalk provide?

- A. A scalable storage appliance on top of Amazon Web Services.
- B. An application container on top of Amazon Web Services.
- C. A service by this name doesn't exist.
- D. A scalable cluster of EC2 instance

Answer: B

NEW QUESTION 231

What will be the status of the snapshot until the snapshot is complete.

- A. running
- B. working
- C. progressing
- D. pending

Answer: D

NEW QUESTION 235

Amazon RDS automated backups and DB Snapshots are currently supported for only the __ storage engine

- A. InnoDB
- B. MyISAM

Answer: A

NEW QUESTION 239

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

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

Answer: B

NEW QUESTION 244

Which is the default region in AWS?

- A. eu-west-1
- B. us-east-1
- C. us-east-2
- D. ap-southeast-1

Answer: B

NEW QUESTION 246

Out of the stripping options available for the EBS volumes, which one has the following disadvantage : 'Doubles the amount of I/O required from the instance to EBS compared to RAID 0, because you're mirroring all writes to a pair of volumes, limiting how much you can stripe.'

- A. Raid 0
- B. RAID 1+0 (RAID 10)
- C. Raid 1
- D. Raid

Answer: B

NEW QUESTION 247

Fill in the blanks: _ let you categorize your EC2 resources in different ways, for example, by purpose, owner, or environment.

- A. wildcards
- B. pointers
- C. Tags
- D. special filters

Answer: C

NEW QUESTION 251

How can I change the security group membership for interfaces owned by other AWS, such as Elastic Load Balancing?

- A. By using the service specific console or API/CLI commands
- B. None of these
- C. Using Amazon EC2 API/CLI
- D. using all these methods

Answer: A

NEW QUESTION 252

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

- A. No
- B. Yes

Answer: A

NEW QUESTION 256

What is Amazon Glacier?

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

Answer: C

NEW QUESTION 260

What is the durability of 53 RRS?

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

Answer: A

NEW QUESTION 264

What happens to the data on an instance if the instance reboots (intentionally or unintentionally)?

- A. Data will be lost
- B. Data persists
- C. Data may persist however cannot be sure

Answer: B

NEW QUESTION 269

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

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

Answer: B

NEW QUESTION 272

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

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

Answer: C

NEW QUESTION 273

Using Amazon IAM, can I give permission based on organizational groups?

- A. Yes but only in certain cases
- B. No
- C. Yes always

Answer: C

NEW QUESTION 275

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

- A. OracleI SAM
- B. MSSQLDB
- C. InnoDB
- D. IV|y|SAIV|

Answer: C

NEW QUESTION 279

When should I choose Provisioned IOPS over Standard RDS storage?

- A. If you have batch-oriented workloads
- B. If you use production online transaction processing (OLTP) workloads.
- C. If you have workloads that are not sensitive to consistent performance

Answer: A

NEW QUESTION 282

What is the minimum charge for the data transferred between Amazon RDS and Amazon EC2 Instances in the same Availability Zone?

- A. USD 0.10 per GB
- B. No charge
- C. It is free.
- D. USD 0.02 per GB
- E. USD 0.01 per GB

Answer: B

NEW QUESTION 285

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

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

Answer: B

NEW QUESTION 289

In regards to IAM you can edit user properties later, but you cannot use the console to change the

- A. user name
- B. password
- C. default group

Answer: A

NEW QUESTION 290

Making your snapshot public shares all snapshot data with everyone. Can the snapshots with AWS Market place product codes be made public?

- A. No

B. Yes

Answer: B

NEW QUESTION 291

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

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

Answer: B

NEW QUESTION 294

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

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

Answer: D

NEW QUESTION 297

Through which of the following interfaces is AWS Identity and Access Management available?

- A) AWS Management Console
- B) Command line interface (CLI)
- C) IAM Query API
- D) Existing libraries

- A. Only through Command line interface (CLI)
- B. A, Band C
- C. A and C
- D. All of the above

Answer: D

NEW QUESTION 299

Select the incorrect statement

- A. In Amazon EC2, the private IP addresses only returned to Amazon EC2 when the instance is stopped or terminated
- B. In Amazon VPC, an instance retains its private IP addresses when the instance is stopped.
- C. In Amazon VPC, an instance does NOT retain its private IP addresses when the instance is stopped.
- D. In Amazon EC2, the private IP address is associated exclusively with the instance for its lifetime

Answer: C

NEW QUESTION 303

How are the EBS snapshots saved on Amazon 53?

- A. Exponentially
- B. Incrementally
- C. EBS snapshots are not stored in the Amazon 53
- D. Decrementally

Answer: B

NEW QUESTION 305

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

- A. TRUE
- B. FALSE

Answer: A

NEW QUESTION 309

What is the maximum response time for a Business level Premium Support case?

- A. 120 seconds
- B. 1 hour
- C. 10 minutes
- D. 12 hours

Answer: B

NEW QUESTION 314

What is the default maximum number of MFA devices in use per AWS account (at the root account level)?

- A. 1
- B. 5
- C. 15
- D. 10

Answer: A

NEW QUESTION 318

Can we attach an EBS volume to more than one EC2 instance at the same time?

- A. Yes.
- B. No
- C. Only EC2-optimized EBS volumes.
- D. Only in read mode.

Answer: A

NEW QUESTION 323

What does Amazon ElastiCache provide?

- A. A service by this name doesn't exist
- B. Perhaps you mean Amazon CloudCache.
- C. A virtual server with a huge amount of memory.
- D. A managed In-memory cache service.
- E. An Amazon EC2 instance with the Memcached software already pre-installed

Answer: C

NEW QUESTION 325

Every user you create in the IAM system starts with __

- A. full permissions
- B. no permissions
- C. partial permissions

Answer: B

NEW QUESTION 328

After an Amazon VPC instance is launched, can I change the VPC security groups it belongs to?

- A. Only if the tag "VPC_Change_Group" is true
- B. Yes
- C. You can.
- D. No
- E. You cannot.
- F. Only if the tag "VPC Change Group" is true

Answer: B

NEW QUESTION 332

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Solution-Architect-Associate Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Solution-Architect-Associate Product From:

<https://www.2passeasy.com/dumps/AWS-Solution-Architect-Associate/>

Money Back Guarantee

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 First Try
- * AWS-Solution-Architect-Associate Practice Test Questions in Multiple Choice Formats and Updates for 1 Year