Amazon

Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate

NEW QUESTION 1
You are trying to launch an EC2 instance, however the instance seems to go into a terminated status immediately. What would probably not be a reason that this is happening?

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

Answer: C

Explanation:
Amazon EC2 provides a virtual computing environments, known as an instance.
After you launch an instance, AWS recommends that you check its status to confirm that it goes from the pending status to the running status, the not terminated status.
The following are a few reasons why an Amazon EBS-backed instance might immediately terminate: You've reached your volume limit.
The AM is missing a required part. The snapshot is corrupt. Reference:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_|nstanceStraightToTerminated.html

NEW QUESTION 2
You have set up an Auto Scaling group. The cool down period for the Auto Scaling group is 7 minutes. The first instance is launched after 3 minutes, while the second instance is launched after 4 minutes. How many minutes after the first instance is launched will Auto Scaling accept another scaling actMty request?

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

Answer: A

Explanation:
If an Auto Scaling group is launching more than one instance, the cool down period for each instance starts after that instance is launched. The group remains locked until the last instance that was launched has completed its cool down period. In this case the cool down period for the first instance starts after 3 minutes and finishes at the 10th minute (3+7 cool down), while for the second instance it starts at the 4th minute and finishes at the 11th minute (4+7 cool down). Thus, the Auto Scaling group will receive another request only after 11 minutes.
Reference:http://docs.aws.amazon.com/AutoScaling/latest/Deve|operGuide/AS_Concepts.htmI

NEW QUESTION 3
Amazon EBS provides the ability to create backups of any Amazon EC2 volume into what is known as

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

Answer: A

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

NEW QUESTION 4
One of the criteria for a new deployment is that the customer wants to use AWS Storage Gateway. However you are not sure whether you should use gateway-cached volumes or gateway-stored volumes or even what the differences are. Which statement below best describes those differences?

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

Answer: A

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

NEW QUESTION 5
A user is storing a large number of objects on AWS S3. The user wants to implement the search functionality among the objects. How can the user achieve this?

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

Answer: D

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

NEW QUESTION 6
After setting up a Virtual Private Cloud (VPC) network, a more experienced cloud engineer suggests that to achieve low network latency and high network throughput you should look into setting up a placement group. You know nothing about this, but begin to do some research about it and are especially curious about its limitations. Which of the below statements is wrong in describing the limitations of a placement group?

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

Answer: B

Explanation:
A placement group is a logical grouping of instances within a single Availability Zone. Using placement groups enables applications to participate in a low-latency, 10 Gbps network. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. To provide the lowest latency, and the highest packet-per-second network performance for your placement group, choose an instance type that supports enhanced networking. Placement groups have the following limitations:
The name you specify for a placement group name must be unique within your AWS account. A placement group can't span multiple Availability Zones.
Although launching multiple instance types into a placement group is possible, this reduces the likelihood that the required capacity will be available for your launch to succeed. We recommend using the same instance type for all instances in a placement group.
You can't merge placement groups. Instead, you must terminate the instances in one placement group, and then relaunch those instances into the other placement group.
A placement group can span peered VPCs; however, you will not get full-bisection bandwidth between instances in peered VPCs. For more information about VPC peering connections, see VPC Peering in the Amazon VPC User Guide.
You can't move an existing instance into a placement group. You can create an AM from your existing instance, and then launch a new instance from the AMI into a placement group.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

NEW QUESTION 7
True orfalsez 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 8
An organization has three separate AWS accounts, one each for development, testing, and production. The organization wants the testing team to have access to certain AWS resources in the production account. How can the organization achieve this?

A. It is not possible to access resources of one account with another account.
B. Create the IAM roles with cross account access.
C. Create the IAM user in a test account, and allow it access to the production environment with the IAM policy.
D. Create the IAM users with cross account acces

Answer: B

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

NEW QUESTION 9
A user wants to use an EBS-backed Amazon EC2 instance for a temporary job. Based on the input data, the job is most likely to finish within a week. Which of the following steps should be followed to terminate the instance automatically once the job is finished?

A. Configure the EC2 instance with a stop instance to terminate it.
B. Configure the EC2 instance with ELB to terminate the instance when it remains idle.
C. Configure the CloudWatch alarm on the instance that should perform the termination action once the instance is idle.
D. Configure the Auto Scaling schedule actMty that terminates the instance after 7 day

Answer: C

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

NEW QUESTION 10
In Amazon EC2, partial instance-hours are billed .

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

Answer: D

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

NEW QUESTION 10
Your EBS volumes do not seem to be performing as expected and your team leader has requested you look into improving their performance. Which of the following is not a true statement relating to the performance of your EBS volumes?

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

Answer: A

Explanation:
Several factors can affect the performance of Amazon EBS volumes, such as instance configuration, I/O characteristics, workload demand, and storage configuration.
Frequent snapshots provide a higher level of data durability, but they may slightly degrade the
performance of your application while the snapshot is in progress. This trade off becomes critical when you have data that changes rapidly. Whenever possible, plan for snapshots to occur during off-peak times in order to minimize workload impact.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSPerformance.htmI

NEW QUESTION 14
A major finance organisation has engaged your company to set up a large data mining application. Using AWS you decide the best service for this is Amazon Elastic MapReduce(EMR) which you know uses Hadoop. Which of the following statements best describes Hadoop?

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

Answer: C

Explanation:
Amazon EMR uses Apache Hadoop as its distributed data processing engine.
Hadoop is an open source, Java software framework that supports data-intensive distributed applications running on large clusters of commodity hardware.
Hadoop implements a programming model named "MapReduce," where the data is dMded into many small fragments of work, each of which may be executed on any node in the cluster.
This framework has been widely used by developers, enterprises and startups and has proven to be a reliable software platform for processing up to petabytes of data on clusters of thousands of commodity machines.
Reference: http://aws.amazon.com/elasticmapreduce/faqs/

NEW QUESTION 17
In Amazon EC2 Container Service, are other container types supported?

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

Answer: C

Explanation:

In Amazon EC2 Container Service, Docker is the only container platform supported by EC2 Container Service presently.
Reference: http://aws.amazon.com/ecs/faqs/

NEW QUESTION 18
In Amazon RDS, security groups are ideally used to:

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

Answer: D

Explanation:
In Amazon RDS, security groups are used to control what IP addresses or EC2 instances can connect to your databases on a DB instance.
When you first create a DB instance, its firewall prevents any database access except through rules specified by an associated security group.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.htmI

NEW QUESTION 23
You need to set up a complex network infrastructure for your organization that will be reasonably easy to deploy, replicate, control, and track changes on. Which AWS service would be best to use to help you accomplish this?

A. AWS Import/Export
B. AWS CIoudFormation
C. Amazon Route 53
D. Amazon CIoudWatch

Answer: B

Explanation:
AWS CIoudFormation is a service that helps you model and set up your Amazon Web Services resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS. You create a template that describes all the AWS resources that you want (like Amazon EC2 instances or Amazon RDS DB instances), and AWS CIoudFormation takes care of provisioning and configuring those resources for you. You don't need to indMdually create and configure AWS resources
and figure out what's dependent on what. AWS CIoudFormation handles all of that.
Reference: http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/WeIcome.htmI

NEW QUESTION 24
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 28
You are setting up a very complex financial services grid and so far it has 5 Elastic IP (EIP) addresses.
You go to assign another EIP address, but all accounts are limited to 5 Elastic IP addresses per region by default, so you aren't able to. What is the reason for this?

A. For security reasons.
B. Hardware restrictions.
C. Public (IPV4) internet addresses are a scarce resource.
D. There are only 5 network interfaces per instanc

Answer: C

Explanation:
Public (IPV4) internet addresses are a scarce resource. There is only a limited amount of public IP space available, and Amazon EC2 is committed to helping use that space efficiently.
By default, all accounts are limited to 5 Elastic IP addresses per region. If you need more than 5 Elastic IP addresses, AWS asks that you apply for your limit to be raised. They will ask you to think through your use case and help them understand your need for additional addresses.
Reference: http://aws.amazon.com/ec2/faqs/#How_many_instances_can_I_run_in_Amazon_EC2

NEW QUESTION 33
Amazon RDS provides high availability and failover support for DB instances using .

A. customized deployments
B. Appstream customizations
C. log events
D. Multi-AZ deployments

Answer: D

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

NEW QUESTION 38
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 42
Having set up a website to automatically be redirected to a backup website if it fails, you realize that there are different types of failovers that are possible. You need all your resources to be available the majority of the time. Using Amazon Route 53 which configuration would best suit this requirement?

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

Answer: A

Explanation:
You can set up a variety of failover configurations using Amazon Route 53 alias: weighted, latency, geolocation routing, and failover resource record sets.
Active-active failover: Use this failover configuration when you want all of your resources to be available the majority of the time. When a resource becomes unavailable, Amazon Route 53 can detect that it's unhealthy and stop including it when responding to queries.
Active-passive failover: Use this failover configuration when you want a primary group of resources to be available the majority of the time and you want a secondary group of resources to be on standby in case all of the primary resources become unavailable. When responding to queries, Amazon Route 53 includes only the healthy primary resources. If all of the primary resources are unhealthy, Amazon Route 53 begins to include only the healthy secondary resources in response to DNS queries.
Active-active-passive and other mixed configurations: You can combine alias and non-alias resource record sets to produce a variety of Amazon Route 53 behaviors.
Reference: http://docs.aws.amazon.com/Route53/latest/DeveIoperGuide/dns-failover.html

NEW QUESTION 45
You decide that you need to create a number of Auto Scaling groups to try and save some money as you have noticed that at certain times most of your EC2 instances are not being used. By default, what is the maximum number of Auto Scaling groups that AWS will allow you to create?

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

Answer: C

Explanation:
Auto Scaling is an AWS service that allows you to increase or decrease the number of EC2 instances within your application's architecture. With Auto Scaling, you create collections of EC2 instances, called Auto Scaling groups. You can create these groups from scratch, or from existing EC2 instances that are already in production.
Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_|imits.htm|#limits_autoscaling

NEW QUESTION 46
A user needs to run a batch process which runs for 10 minutes. This will only be run once, or at maximum twice, in the next month, so the processes will be temporary only. The process needs 15 X-Large instances. The process downloads the code from S3 on each instance when it is launched, and then generates a temporary log file. Once the instance is terminated, all the data will be lost. Which of the below mentioned pricing models should the user choose in this case?

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

Answer: A

Explanation:
In Amazon Web Services, the spot instance is useful when the user wants to run a process temporarily. The spot instance can terminate the instance if the other user outbids the existing bid. In this case all storage is temporary and the data is not required to be persistent. Thus, the spot instance is a good option to save money.
Reference: http://aws.amazon.com/ec2/purchasing-options/spot-instances/

NEW QUESTION 51
You have been storing massive amounts of data on Amazon Glacier for the past 2 years and now start to wonder if there are any limitations on this. What is the correct answer to your QUESTION ?

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

Answer: C

Explanation:
An archive is a durably stored block of information. You store your data in Amazon Glacier as archives. You may upload a single file as an archive, but your costs will be lower if you aggregate your data. TAR and ZIP are common formats that customers use to aggregate multiple files into a single file before uploading to Amazon Glacier.
The total volume of data and number of archives you can store are unlimited. IndMdual Amazon Glacier archives can range in size from 1 byte to 40 terabytes.
The largest archive that can be uploaded in a single upload request is 4 gigabytes.
For items larger than 100 megabytes, customers should consider using the MuItipart upload capability. Archives stored in Amazon Glacier are immutable, i.e. archives can be uploaded and deleted but cannot be edited or overwritten.
Reference: https://aws.amazon.com/gIacier/faqs/

NEW QUESTION 55
A user has configured ELB with two EBS backed EC2 instances. The user is trying to understand the DNS access and IP support for ELB. Which of the below mentioned statements may not help the user understand the IP mechanism supported by ELB?

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

Answer: D

Explanation:
Elastic Load Balancing supports both Internet Protocol version 6 (IPv6) and Internet Protocol version 4 (IPv4). Clients can connect to the user's load balancer using either IPv4 or IPv6 (in EC2-Classic) DNS. However, communication between the load balancer and its back-end instances uses only IPv4. The user can use the Dualstack-prefixed DNS name to enable IPv6 support for communications between the client and the load balancers. Thus, the clients are able to access the load balancer using either IPv4 or IPv6 as their indMdual connectMty needs dictate.
Reference: http://docs.aws.amazon.com/EIasticLoadBaIancing/latest/DeveIoperGuide/UserScenariosForEC2.html

NEW QUESTION 60
You need to set up a high level of security for an Amazon Relational Database Service (RDS) you have just built in order to protect the confidential information stored in it. What are all the possible security groups that RDS uses?

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

Answer: A

Explanation:
A security group controls the access to a DB instance. It does so by allowing access to IP address ranges or Amazon EC2 instances that you specify.
Amazon RDS uses DB security groups, VPC security groups, and EC2 security groups. In simple terms, a DB security group controls access to a DB instance that is not in a VPC, a VPC security group controls access to a DB instance inside a VPC, and an Amazon EC2 security group controls access to an EC2 instance and can be used with a DB instance.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html

NEW QUESTION 61
You have been using T2 instances as your CPU requirements have not been that intensive. However you now start to think about larger instance types and start looking at M and IV|3 instances. You are a little confused as to the differences between them as they both seem to have the same ratio of CPU and memory. Which statement below is incorrect as to why you would use one over the other?

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

Answer: B

Explanation:

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

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

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

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

However, if you need more disk storage than what is provided in M3 instances, you may still find M1 instances useful for running your applications.
Reference: https://aws.amazon.com/ec2/faqs/

NEW QUESTION 65

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

A. Mu|ti AZ
B. Read Replica
C. Multi region
D. PostgreSQL does not support HA

Answer: A

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

NEW QUESTION 66

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 devic

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 67

A scope has been handed to you to set up a super fast gaming server and you decide that you will use Amazon DynamoDB as your database. For efficient access to data in a table, Amazon DynamoDB creates and maintains indexes for the primary key attributes. A secondary index is a data structure that contains a subset of attributes from a table, along with an alternate key to support Query operations. How many types of secondary indexes does DynamoDB support?

A. 2
B. 16
C. 4
D. As many as you nee

Answer: A

Explanation:
DynamoDB supports two types of secondary indexes:
Local secondary index — an index that has the same hash key as the table, but a different range key. A local secondary index is "local" in the sense that every partition of a local secondary index is scoped to a table partition that has the same hash key.
Global secondary index — an index with a hash and range key that can be different from those on the table. A global secondary index is considered "global" because queries on the index can span all of the data in a table, across all partitions.
Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html

NEW QUESTION 69

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

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

Answer: D

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

NEW QUESTION 72

A user is planning to make a mobile game which can be played online or offline and will be hosted on EC2.
The user wants to ensure that if someone breaks the highest score or they achieve some milestone they can inform all their colleagues through email. Which of the below mentioned AWS services helps achieve this goal?

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

Answer: B

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

NEW QUESTION 77
You have multiple VPN connections and want to provide secure communication between sites using the AWS VPN CIoudHub. Which statement is the most accurate in describing what you must do to set this up correctly?

A. Create a virtual private gateway with multiple customer gateways, each with unique Border Gateway Protocol (BGP) Autonomous System Numbers (ASNs)
B. Create a virtual private gateway with multiple customer gateways, each with a unique set of keys
C. Create a virtual public gateway with multiple customer gateways, each with a unique Private subnet
D. Create a virtual private gateway with multiple customer gateways, each with unique subnet id

Answer: A

Explanation:
If you have multiple VPN connections, you can provide secure communication between sites using the AWS VPN CIoudHub. The VPN CIoudHub operates on a simple hub-and-spoke model that you can use with or without a VPC. This design is suitable for customers with multiple branch offices and existing Internet connections who'd like to implement a convenient, potentially low-cost hub-and-spoke model for primary or backup connectMty between these remote offices.
To use the AWS VPN CIoudHub, you must create a virtual private gateway with multiple customer
gateways, each with unique Border Gateway Protocol (BGP) Autonomous System Numbers (ASNs). Customer gateways advertise the appropriate routes (BGP prefixes) over their VPN connections. These routing advertisements are received and re-advertised to each BGP peer, enabling each site to send data to and receive data from the other sites. The routes for each spoke must have unique ASNs and the sites must not have overlapping IP ranges. Each site can also send and receive data from the VPC as if they were using a standard VPN connection.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPN_CIoudHub.htmI

NEW QUESTION 81
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 grou

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.htm|

NEW QUESTION 86
You are configuring a new VPC for one of your clients for a cloud migration project, and only a public VPN will be in place. After you created your VPC, you created a new subnet, a new internet gateway, and attached your internet gateway to your VPC. When you launched your first instance into your VPC, you realized that you aren't able to connect to the instance, even if it is configured with an elastic IP. What should be done to access the instance?

A. A route should be created as 0.0.0.0/0 and your internet gateway as target.
B. Attach another ENI to the instance and connect via new ENI.
C. A NAT instance should be created and all traffic should be forwarded to NAT instance.
D. A NACL should be created that allows all outbound traffi

Answer: A

Explanation:
All traffic should be routed via Internet Gateway. So, a route should be created with 0.0.0.0/0 as a source, and your Internet Gateway as your target.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario1.htmI

NEW QUESTION 89
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_To|erant_Applications.pdf

NEW QUESTION 93
Doug has created a VPC with CIDR 10.201.0.0/16 in his AWS account. In this VPC he has created a public subnet with CIDR block 10.201.31.0/24. While launching a new EC2 from the console, he is not able to assign the private IP address 10.201.31.6 to this instance. Which is the most likely reason for this issue?

A. Private IP address 10.201.31.6 is blocked via ACLs in Amazon infrastructure as a part of platform security.
B. Private address IP 10.201.31.6 is currently assigned to another interface.
C. Private IP address 10.201.31.6 is not part of the associated subnet's IP address range.
D. Private IP address 10.201.31.6 is reserved by Amazon for IP networking purpose

Answer: B

Explanation:
In Amazon VPC, you can assign any Private IP address to your instance as long as it is: Part of the associated subnet's IP address range
Not reserved by Amazon for IP networking purposes Not currently assigned to another interface Reference: http://aws.amazon.com/vpc/faqs/

NEW QUESTION 96
You need to create a JSON-formatted text file for AWS CIoudFormation. This is your first template and the only thing you know is that the templates include several major sections but there is only one that is required for it to work. What is the only section required?

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

Answer: C

Explanation:
AWS CIoudFormation is a service that helps you model and set up your Amazon Web Services resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS. You create a template that describes all the AWS resources that you want (like Amazon EC2 instances or Amazon RDS DB instances), and AWS CIoudFormation takes care of provisioning and configuring those resources for you.
A template is a JSON-formatted text file that describes your AWS infrastructure. Templates include several major sections.
The Resources section is the only section that is required.
The first character in the template must be an open brace ({), and the last character must be a closed brace (}). The following template fragment shows the template structure and sections.
Reference: http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/template-anatomy.html

NEW QUESTION 98
You are planning and configuring some EBS volumes for an application. In order to get the most performance out of your EBS volumes, you should attach them to an instance with enough to support your volumes.

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

Answer: C

Explanation:
When you plan and configure EBS volumes for your application, it is important to consider the configuration of the instances that you will attach the volumes to. In order to get the most performance out of your EBS volumes, you should attach them to an instance with enough bandwidth to support your volumes, such as an EBS-optimized instance or an instance with 10 Gigabit network connectMty. This is especially important when you use General Purpose (SSD) or Provisioned IOPS (SSD) volumes, or when you stripe multiple volumes together in a RAID configuration.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-ec2-config.htmI

NEW QUESTION 100
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 102
A user is hosting a website in the US West-1 region. The website has the highest client base from the Asia-Pacific (Singapore / Japan) region. The application is accessing data from S3 before serving it to client. Which of the below mentioned regions gives a better performance for S3 objects?

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

Answer: D

Explanation:
Access to Amazon S3 from within Amazon EC2 in the same region is fast. In this aspect, though the client base is Singapore, the application is being hosted in the US West-1 region. Thus, it is recommended that S3 objects be stored in the US-West-1 region.
Reference: http://media.amazonwebservices.com/AWS_Storage_Options.pdf

NEW QUESTION 105
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 resourc

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 107
You have been asked to tighten up the password policies in your organization after a serious security breach, so you need to consider every possible security measure. Which of the following is not an account password policy for IAM Users that can be set?

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

Answer: C

Explanation:
IAM users need passwords in order to access the AWS Management Console. (They do not need passwords if they will access AWS resources programmatically by using the CLI, AWS SDKs, or the APIs.)
You can use a password policy to do these things: Set a minimum password length.
Require specific character types, including uppercase letters, lowercase letters, numbers, and non-alphanumeric characters. Be sure to remind your users that passwords are case sensitive. Allow all IAM users to change their own passwords.
Require IAM users to change their password after a specified period of time (enable password expiration). Prevent IAM users from reusing previous passwords.
Force IAM users to contact an account administrator when the user has allowed his or her password to expue.
Reference: http://docs.aws.amazon.com/|AM/latest/UserGuide/Using_ManagingPasswordPolicies.htm|

NEW QUESTION 110
Your organization is in the business of architecting complex transactional databases. For a variety of reasons, this has been done on EBS. What is AWS's recommendation for customers who have architected databases using EBS for backups?

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

Answer: A

Explanation:
Data stored in Amazon EBS volumes is redundantly stored in multiple physical locations as part of normal operation of those services and at no additional charge. However, Amazon EBS replication is stored within the same availability zone, not across multiple zones; therefore, it is highly recommended that you conduct regular snapshots to Amazon S3 for long-term data durability.
For customers who have architected complex transactional databases using EBS, it is recommended that backups to Amazon S3 be performed through the database management system so that distributed transactions and logs can be checkpointed.
AWS does not perform backups of data that are maintained on virtual disks attached to running instances on Amazon EC2.
Reference: http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf

NEW QUESTION 112
You have three Amazon EC2 instances with Elastic IP addresses in the US East (Virginia) region, and you want to distribute requests across all three IPs evenly for users for whom US East (Virginia) is the appropriate region.
How many EC2 instances would be sufficient to distribute requests in other regions?

A. 3
B. 9

C. 2
D. 1

Answer: D

Explanation:
If your application is running on Amazon EC2 instances in two or more Amazon EC2 regions, and if you have more than one Amazon EC2 instance in one or more regions, you can use latency-based routing to route traffic to the correct region and then use weighted resource record sets to route traffic to instances within the region based on weights that you specify.
For example, suppose you have three Amazon EC2 instances with Elastic IP addresses in the US East (Virginia) region and you want to distribute requests across all three IPs evenly for users for whom US East (Virginia) is the appropriate region. Just one Amazon EC2 instance is sufficient in the other regions, although you can apply the same technique to many regions at once.
Reference: http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/Tutorials.html

NEW QUESTION 117
A major client who has been spending a lot of money on his internet service provider asks you to set up an AWS Direct Connection to try and save him some money. You know he needs high-speed connectMty. Which connection port speeds are available on AWS Direct Connect?

A. 500Mbps and 1Gbps
B. 1Gbps and 10Gbps
C. 100Mbps and 1Gbps
D. 1Gbps

Answer: B

Explanation:
AWS Direct Connect is a network service that provides an alternative to using the internet to utilize AWS cloud services.
Using AWS Direct Connect, data that would have previously been transported over the Internet can now be delivered through a private network connection between AWS and your datacenter or corporate network.
1Gbps and 10Gbps ports are available. Speeds of 50Mbps, 100Mbps, 200Mbps, 300Mbps, 400Mbps, and 500Mbps can be ordered from any APN partners supporting AWS Direct Connect.
Reference: https://aws.amazon.com/directconnect/faqs/

NEW QUESTION 120
After setting up some EC2 instances you now need to set up a monitoring solution to keep track of these instances and to send you an email when the CPU hits a certain threshold. Which statement below best describes what thresholds you can set to trigger a CIoudWatch Alarm?

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

Answer: A

Explanation:
Amazon CIoudWatch is a monitoring service for AWS cloud resources and the applications you run on AWS. You can use Amazon CIoudWatch to collect and track metrics, collect and monitor log files, and set
alarms.
When you create an alarm, you first choose the Amazon CIoudWatch metric you want it to monitor. Next, you choose the evaluation period (e.g., five minutes or one hour) and a statistical value to measure (e.g., Average or Maximum).
To set a threshold, set a target value and choose whether the alarm will trigger when the value is greater than (>), greater than or equal to (>=), less than (<), or less than or equal to (<=) that value.
Reference: http://aws.amazon.com/cIoudwatch/faqs/

NEW QUESTION 122
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 SimpIeDB

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 scaHng.
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 127
A favored client needs you to quickly deploy a database that is a relational database service with minimal administration as he wants to spend the least amount of time administering it. Which database would be the best option?

A. Amazon Simp|eDB
B. Your choice of relational AMs on Amazon EC2 and EBS.
C. Amazon RDS
D. Amazon Redshift

Answer: C

Explanation:
Amazon Relational Database Service (Amazon RDS) is a web service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while managing time-consuming database administration tasks, freeing you up to focus on your applications and business.
Amazon RDS gives you access to the capabilities of a familiar MySQL, Oracle, SQL Server, or PostgreSQL database engine. This means that the code, applications, and tools you already use today with your existing databases can be used with Amazon RDS. Amazon RDS automatically patches the database software and backs up your database, storing the backups for a user-defined retention period and enabling point-in-time recovery.
Reference: https://aws.amazon.com/running_databases/#rds_anchor

NEW QUESTION 128
You're trying to delete an SSL certificate from the IAM certificate store, and you're getting the message "Certificate: <certificate-id> is being used by CloudFront." Which of the following statements is probably the reason why you are getting this error?

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

Answer: A

Explanation:
CloudFront is a web service that speeds up distribution of your static and dynamic web content, for example, .html, .css, .php, and image files, to end users. Every CloudFront web distribution must be associated either with the default CloudFront certificate or with a custom SSL certificate. Before you can delete an SSL certificate, you need to either rotate SSL certificates (replace the current custom SSL certificate with another custom SSL certificate) or revert from using a custom SSL certificate to using the default CloudFront certificate.
Reference: http://docs.aws.amazon.com/AmazonCloudFront/latest/Deve|operGuide/Troubleshooting.htm|

NEW QUESTION 131
In relation to AWS CloudHSM, High-availability (HA) recovery is hands-off resumption by failed HA group members.
Prior to the introduction of this function, the HA feature provided redundancy and performance, but required that a failed/lost group member be reinstated.

A. automatically
B. periodically
C. manually
D. continuosly

Answer: C

Explanation:
In relation to AWS CloudHS|VI, High-availability (HA) recovery is hands-off resumption by failed HA group members.
Prior to the introduction of this function, the HA feature provided redundancy and performance, but required that a failed/lost group member be manually reinstated.
Reference: http://docs.aws.amazon.com/cloudhsm/latest/userguide/ha-best-practices.html

NEW QUESTION 133
Amazon Elastic Load Balancing is used to manage traffic on a fileet 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 fileet 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 indMdual instance. Reference:
http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf

NEW QUESTION 138
You have been asked to build AWS infrastructure for disaster recovery for your local applications and within that you should use an AWS Storage Gateway as part of the solution. Which of the following best describes the function of an AWS Storage Gateway?

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

Answer: C

Explanation:
AWS Storage Gateway connects an on-premises software appliance with cloud-based storage to provide seamless integration with data security features between your on-premises IT environment and the Amazon Web Services (AWS) storage infrastructure. You can use the service to store data in the AWS cloud for scalable and cost-effective storage that helps maintain data security. AWS Storage Gateway offers both volume-based and tape-based storage solutions:
Volume gateways Gateway-cached volumes Gateway-stored volumes
Gateway-virtual tape library (VTL)
Reference:
http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_disasterrecovery_07.pdf

NEW QUESTION 143
A government client needs you to set up secure cryptographic key storage for some of their extremely confidential data. You decide that the AWS CIoudHSM is the best service for this. However, there seem to be a few pre-requisites before this can happen, one of those being a security group that has certain ports open. Which of the following is correct in regards to those security groups?

A. A security group that has port 22 (for SSH) or port 3389 (for RDP) open to your network.
B. A security group that has no ports open to your network.
C. A security group that has only port 3389 (for RDP) open to your network.
D. A security group that has only port 22 (for SSH) open to your network.

Answer: A

Explanation:
AWS CIoudHSM provides secure cryptographic key storage to customers by making hardware security modules (HSMs) available in the AWS cloud.
AWS C|oudHSM requires the following environment before an HSM appliance can be provisioned. A virtual private cloud (VPC) in the region where you want the AWS CIoudHSM service.
One private subnet (a subnet with no Internet gateway) in the VPC. The HSM appliance is provisioned into this subnet.
One public subnet (a subnet with an Internet gateway attached). The control instances are attached to this subnet.
An AWS Identity and Access Management (IAM) role that delegates access to your AWS resources to AWS CIoudHSM.
An EC2 instance, in the same VPC as the HSM appliance, that has the SafeNet client software installed. This instance is referred to as the control instance and is used to connect to and manage the HSM appliance.
A security group that has port 22 (for SSH) or port 3389 (for RDP) open to your network. This security group is attached to your control instances so you can access them remotely.

NEW QUESTION 145
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 147
You are setting up your first Amazon Virtual Private Cloud (Amazon VPC) network so you decide you should probably use the AWS Management Console and the VPC Wizard. Which of the following is not an option for network architectures after launching the "Start VPC Wizard" in Amazon VPC page on the AWS Management Console?

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

Answer: B

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

VPC to add more subnets or add or remove gateways at any time after the VPC has been created.
The four options are:
VPC with a Single Public Subnet Only VPC with Public and Private Subnets
VPC with Public and Private Subnets and Hardware VPN Access VPC with a Private Subnet Only and Hardware VPN Access Reference:
https://aws.amazon.com/vpc/faqs/

NEW QUESTION 150
Which one of the below doesn't affect Amazon CloudFront billing?

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

Answer: A

Explanation:
Amazon CloudFront is a web service for content delivery. CloudFront delivers your content using a global network of edge locations and works seamlessly with Amazon S3 which durably stores the original and definitive versions of your files.
Amazon CloudFront billing is maily affected by Data Transfer Out
Edge Location Traffic Distribution Requests
Dedicated IP SSL Certificates
Reference: http://calcu|ator.s3.amazonaws.com/index.html

NEW QUESTION 151
A user has defined an AutoScaling termination policy to first delete the instance with the nearest billing hour. AutoScaling has launched 3 instances in the US-East-1A region and 2 instances in the US-East-1 B region. One of the instances in the US-East-1B region is running nearest to the billing hour. Which instance will AutoScaling terminate first while executing the termination action?

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

Answer: C

Explanation:
Even though the user has configured the termination policy, before AutoScaling selects an instance to terminate, it first identifies the Availability Zone that has more instances than the other Availability Zones used by the group. Within the selected Availability Zone, it identifies the instance that matches the specified termination policy.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/us-termination-policy.html

NEW QUESTION 152
You are playing around with setting up stacks using JSON templates in CloudFormation to try and understand them a little better. You have set up about 5 or 6 but now start to wonder if you are being charged for these stacks. What is AWS's billing policy regarding stack resources?

A. You are not charged for the stack resources if they are not taking any traffic.
B. You are charged for the stack resources for the time they were operating (even if you deleted the stack right away)
C. You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 60 minutes)
D. You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 30 minutes)

Answer: B

Explanation:
A stack is a collection of AWS resources that you can manage as a single unit. In other words, you can create, update, or delete a collection of resources by creating, updating, or deleting stacks. All the resources in a stack are defined by the stack's AWS CloudFormation template. A stack, for instance, can include all the resources required to run a web application, such as a web server, a database, and networking rules. If you no longer require that web application, you can simply delete the stack, and all of its related resources are deleted.
You are charged for the stack resources for the time they were operating (even if you deleted the stack right away).
Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacks.html

NEW QUESTION 155
You need to create a management network using network interfaces for a virtual private cloud (VPC) network. Which of the following statements is incorrect pertaining to Best Practices for Configuring Network Interfaces.

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

Answer: D

Explanation:
Best Practices for Configuring Network Interfaces
You can attach a network interface to an instance when it's running (hot attach), when it's stopped (warm attach), or when the instance is being launched (cold attach).

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

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

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

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

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

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

Reference:

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

NEW QUESTION 160
Can I change the EC2 security groups after an instance is launched in EC2-Classic?

A. Yes, you can change security groups after you launch an instance in EC2-Classic.
B. No, you cannot change security groups after you launch an instance in EC2-Classic.
C. Yes, you can only when you remove rules from a security group.
D. Yes, you can only when you add rules to a security grou

Answer: B

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.htmI

NEW QUESTION 164
You can seamlessly join an EC2 instance to your directory domain. What connectMty do you need to be able to connect remotely to this instance?

A. You must have IP connectMty to the instance from the network you are connecting from.
B. You must have the correct encryption keys to connect to the instance remotely.
C. You must have enough bandwidth to connect to the instance.
D. You must use MFA authentication to be able to connect to the instance remotel

Answer: A

Explanation:
You can seamlessly join an EC2 instance to your directory domain when the instance is launched using the Amazon EC2 Simple Systems Manager. If you need to manuallyjoin an EC2 instance to your domain, you must launch the instance in the proper region and security group or subnet, then join the instance to the domain. To be able to connect remotely to these instances, you must have IP connectMty to the instances from the network you are connecting from. In most cases, this requires that an Internet gateway be attached to your VPC and that the instance has a public IP address.
Reference: http://docs.aws.amazon.com/directoryservice/latest/admin-guide/join_a_directory.html

NEW QUESTION 166
In Amazon EC2, you are billed instance-hours when .

A. your EC2 instance is in a running state
B. the instance exits from Amazon S3 console
C. your instance still exits the EC2 console
D. EC2 instances stop

Answer: A

Explanation:
You are billed instance-hours as long as your EC2 instance is in a running state. Reference: http://aws.amazon.com/ec2/faqs/

NEW QUESTION 169
A user is making a scalable web application with compartmentalization. The user wants the log module to be able to be accessed by all the application functionalities in an asynchronous way. Each module of the application sends data to the log module, and based on the resource availability it will process the logs. Which AWS service helps this functionality?

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

Answer: A

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

NEW QUESTION 173
How can you apply more than 100 rules to an Amazon EC2-Classic?

A. By adding more security groups
B. You need to create a default security group specifying your required rules if you need to use more than 100 rules per security group.
C. By default the Amazon EC2 security groups support 500 rules.
D. You can't add more than 100 rules to security groups for an Amazon EC2 instanc

Answer: D

Explanation:
In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 100 rules to a security group.
Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-network-security.htmI

NEW QUESTION 177
Your company has HQ in Tokyo and branch offices all over the world and is using a logistics software with a multi-regional deployment on AWS in Japan, Europe and USA, The logistic software has a 3- tier architecture and currently uses MySQL 5.6 for data persistence. Each region has deployed its own database
In the HQ region you run an hourly batch process reading data from every region to compute cross regional reports that are sent by email to all offices this batch process must be completed as fast as possible to quickly optimize logistics how do you build the database architecture in order to meet the requirements'?

A. For each regional deployment, use RDS MySQL with a master in the region and a read replica in the HQ region
B. For each regional deployment, use MySQL on EC2 with a master in the region and send hourly EBS snapshots to the HQ region
C. For each regional deployment, use RDS MySQL with a master in the region and send hourly RDS snapshots to the HQ region
D. For each regional deployment, use MySQL on EC2 with a master in the region and use 53 to copy data files hourly to the HQ region
E. Use Direct Connect to connect all regional MySQL deployments to the HQ region and reduce network latency for the batch process

Answer: A

NEW QUESTION 179
You have launched an EC2 instance with four (4) 500GB EBS Provisioned IOPS volumes attached The EC2 Instance Is EBS-Optimized and supports 500 Mbps throughput between EC2 and EBS The two EBS volumes are configured as a single RAID o 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 1/0 performance of the instance, you add an additional two 500 GB EBS Provisioned IOPS volumes to the RAID Each volume Is provisioned to 4.000 IOPs like the original four for a total of 24.000 IOPS on the EC2 instance Monitoring shows that the EC2 instance CPU utilization increased from 50% to 70%. but the total random IOPS measured at the instance level does not increase at all.
What is the problem and a valid solution?

A. Larger storage volumes support higher Provisioned IOPS rates: increase the provisioned volumestorage of each of the 6 EBS volumes to ITB
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 1'0 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 instant root volume to also be a 500GB 4.000 Provisioned IOPS volume.

Answer: E

NEW QUESTION 182
You need a persistent and durable storage to trace call actMty of an IVR (Interactive Voice Response) system. Call duration is mostly in the 2-3 minutes timeframe. Each traced call can be either active or terminated. An external application needs to know each minute the list of currently active calls, which are usually a few calls/second. Put once per month there is a periodic peak up to 1000 calls/second for a few hours. The system is open 24/7 and any downtime should be avoided.
Historical data is periodically archived to files. Cost saving is a priority for this project.
What database implementation would better fit this scenario, keeping costs as low as possible?

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

Answer: A

NEW QUESTION 186
Your application is using an ELB in front of an Auto Scaling group of web/application sewers deployed across two AZs and a MuIti-AZ RDS Instance for data persistence.
The database CPU is often above 80% usage and 90% of 1/0 operations on the database are reads. To improve performance you recently added a single-node Memcached ElastiCache Cluster to cache frequent DB query results. In the next weeks the overall workload is expected to grow by 30%.
Do you need to change anything in the architecture to maintain the high availability or the application with the anticipated additional load? Why?

A. Yes, you should deploy two Memcached ElastiCache Clusters in different AZs because the RDS instance will not be able to handle the load if the cache node fails.
B. No, if the cache node fails you can always get the same data from the DB without having any availability impact.
C. No, if the cache node fails the automated ElastiCache node recovery feature will prevent any availability impact.
D. Yes, you should deploy the Memcached ElastiCache Cluster with two nodes in the same AZ as the RDS DB master instance to handle the load if one cache node fails.

Answer: A

Explanation:
ElastiCache for Memcached
The primary goal of caching is typically to offload reads from your database or other primary data source. In most apps, you have hot spots of data that are regularly queried, but only updated periodically. Think of the front page of a blog or news site, or the top 100 leaderboard in an online game. In this type of case, your app can receive dozens, hundreds, or even thousands of requests for the same data before it's updated again. Having your caching layer handle these queries has several advantages. First, it's considerably cheaper to add an in-memory cache than to scale up to a larger database cluster. Second, an in-memory cache is also easier to scale out, because it's easier to distribute an in-memory cache horizontally than a relational database.
Last, a caching layer provides a request buffer in the event of a sudden spike in usage. If your app or game ends up on the front page of Reddit or the App Store, it's not unheard of to see a spike that is 10 to 100 times your normal application load. Even if you autoscale your application instances, a IOx request spike will likely make your database very unhappy.
Let's focus on ElastiCache for Memcached first, because it is the best fit for a caching focused solution. We'll revisit Redislater in the paper, and weigh its advantages and disadvantages.
Architecture with ElastiCache for Memcached
When you deploy an ElastiCache Memcached cluster, it sits in your application as a separate tier alongside your database. As mentioned previously, Amazon ElastiCache does not directly communicate with your database tier, or indeed have any particular knowledge of your database. A simplified deployment for a web application looks something like this:

In this architecture diagram, the Amazon EC2 application instances are in an Auto Scaling group, located behind a load balancer using Elastic Load Balancing, which distributes requests among the instances. As requests come into a given EC2 instance, that EC2 instance is responsible for communicating with ElastiCache and the database tier. For development purposes, you can begin with a single ElastiCache node to test your application, and then scale to additional cluster nodes by modifying t he ElastiCache cluster. As you add additional cache nodes, the EC2 application instances are able to distribute cache keys across multiple ElastiCache nodes. The most common practice is to use client-side sharding to distribute keys across cache nodes, which we will discuss later in this paper.

When you launch an ElastiCache cluster, you can choose the Availability Zone(s) that the cluster lives in. For best performance, you should configure your cluster to use the same Availability Zones as your application servers. To launch an ElastiCache cluster in a specific Availability Zone, make sure to specify the Preferred Zone(s) option during cache cluster creation. The Availability Zones that you specify will be where ElastiCache will launch your cache nodes. We recommend that you select Spread Nodes Across Zones, which tells ElastiCache to distribute cache nodes across these zones as evenly as possible. This distribution will mitigate the impact of an Availability Zone disruption on your E|astiCache nodes. The trade-off is that some of the requests from your application to ElastiCache will go to a node in a different Availability Zone, meaning latency will be slightly higher.
For more details, refer to Creating a Cache Cluster in the Amazon ElastiCache User Guide.
As mentioned at the outset, ElastiCache can be coupled with a wide variety of databases. Here is an example architecture that uses Amazon DynamoDB instead of Amazon RDS and IV|ySQL:

This combination of DynamoDB and ElastiCache is very popular with mobile and game companies, because DynamoDB allows for higher write throughput at lower cost than traditional relational databases. In addition, DynamoDB uses a key-value access pattern similar to ElastiCache, which also simplifies the programming model. Instead of using relational SQL for the primary database but then key-value patterns for the cache, both the primary database and cache can be programmed similarly.
In this architecture pattern, DynamoDB remains the source of truth for data, but application reads are offloaded to ElastiCache for a speed boost.

NEW QUESTION 189
Your startup wants to implement an order fulfillment process for selling a personalized gadget that needs an average of 3-4 days to produce with some orders taking up to 6 months you expect 10 orders per day on your first day. 1000 orders per day after 6 months and 10,000 orders after 12 months.
Orders coming in are checked for consistency men dispatched to your manufacturing plant for production quality control packaging shipment and payment processing If the product does not meet the quality standards at any stage of the process employees may force the process to repeat a step Customers are notified via email about order status and any critical issues with their orders such as payment failure.
Your case architecture includes AWS Elastic Beanstalk for your website with an RDS MySQL instance for customer data and orders.
How can you implement the order fulfillment process while making sure that the emails are delivered reliably?

A. Add a business process management application to your Elastic Beanstalk app servers and re-use the ROS database for tracking order status use one of the Elastic Beanstalk instances to send emails to customers.
B. Use SWF with an Auto Scaling group of actMty workers and a decider instance in another Auto Scaling group with min/max=I Use the decider instance to send emails to customers.
C. Use SWF with an Auto Scaling group of actMty workers and a decider instance in another Auto Scaling group with min/max=I use SES to send emails to customers.
D. Use an SOS queue to manage all process tasks Use an Auto Scaling group of EC2 Instances that poll the tasks and execute the
E. Use SES to send emails to customers.

Answer: C

NEW QUESTION 192
Your company previously configured a heavily used, dynamically routed VPN connection between your on-premises data center and AWS. You recently provisioned a DirectConnect connection and would like to start using the new connection. After configuring DirectConnect settings in the AWS Console, which of

the following options win provide the most seamless transition for your users?

A. Delete your existing VPN connection to avoid routing loops configure your DirectConnect router with the appropriate settings and verity network traffic is leveraging DirectConnect.
B. Configure your DirectConnect router with a higher 8GP priority man your VPN router, verify network traffic is leveraging Directconnect and then delete your existing VPN connection.
C. Update your VPC route tables to point to the DirectConnect connection configure your DirectConnect router with the appropriate settings verify network traffic is leveraging DirectConnect and then delete the VPN connection.
D. Configure your DirectConnect router, update your VPC route tables to point to the DirectConnect connection, configure your VPN connection with a higher BGP point
E. And verify network traffic is leveraging the DirectConnect connection.

Answer: D

NEW QUESTION 196
You are implementing AWS Direct Connect. You intend to use AWS public service end points such as Amazon 53, across the AWS Direct Connect link. You want other Internet traffic to use your existing link to an Internet Service Provider.
What is the correct way to configure AW5 Direct connect for access to services such as Amazon 53?

A. Configure a public Interface on your AW5 Direct Connect link Configure a static route via your AW5 Direct Connect link that points to Amazon 53 Advertise a default route to AW5 using BGP.
B. Create a private interface on your AW5 Direct Connect lin
C. Configure a static route via your AW5 Direct connect link that points to Amazon 53 Configure specific routes to your network in your VPC,
D. Create a public interface on your AW5 Direct Connect link Redistribute BGP routes into your existing routing infrastructure advertise specific routes for your network to AW5.
E. Create a private interface on your AW5 Direct connect lin
F. Redistribute BGP routes into your existing routing infrastructure and advertise a default route to AW5.

Answer: C

NEW QUESTION 198
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 prMlege and there must be controls in place to ensure that the credentials used by the 5aa5 vendor cannot be used by any other third party. Which of the following would meet all of these conditions?

A. From the AW5 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 tor 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:Get0bject-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 ro Ie 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 200
You are designing a data leak prevention solution for your VPC environment. You want your VPC Instances to be able to access software depots and distributions on the Internet for product updates. The depots and distributions are accessible via third party CONs by their URLs. You want to explicitly deny any other outbound connections from your VPC instances to hosts on the internet.
Which of the following options would you consider?

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

Answer: A

NEW QUESTION 204
An administrator is using Amazon CIoudFormation to deploy a three tier web application that consists of a web tier and application tier that will utilize Amazon DynamoDB for storage when creating the CIoudFormation template which of the following would allow the application instance access to the DynamoDB tables without exposing API credentials?

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

Answer: C

NEW QUESTION 205
An AWS customer is deploying an application mat is composed of an AutoScaling group of EC2 Instances.
The customers security policy requires that every outbound connection from these instances to any other service within the customers Virtual Private Cloud must be authenticated using a unique x 509 certificate that contains the specific instance-id.
In addition an x 509 certificates must Designed by the customer's Key management service in order to be trusted for authentication.
Which of the following configurations will support these requirements?

A. Configure an IAM Role that grants access to an Amazon 53 object containing a signed certificate and configure me Auto Scaling group to launch instances with this role Have the instances bootstrap get the certificate from Amazon 53 upon first boot.
B. Embed a certificate into the Amazon Machine Image that is used by the Auto Scaling group Have the launched instances generate a certificate signature request with the instance's assigned instance- id to the Key management service for signature.
C. Configure the Auto Scaling group to send an SNS notification of the launch of a new instance to the trusted key management servic
D. Have the Key management service generate a signed certificate and send it directly to the newly launched instance.
E. Configure the launched instances to generate a new certificate upon first boot Have the Key management service poll the AutoScaling group for associated instances and send new instances acertificate signature (hat contains the specific instance-i

Answer: A

NEW QUESTION 210
You are designing an intrusion detection prevention (IDS/IPS) solution for a customer web application in a single VPC. You are considering the options for implementing IOS IPS protection for traffic coming from the Internet.
Which of the following options would you consider? (Choose 2 answers)

A. Implement IDS/IPS agents on each Instance running In VPC
B. Configure an instance in each subnet to switch its network interface card to promiscuous mode and analyze network traffic.
C. Implement Elastic Load Balancing with SSL listeners In front of the web applications
D. Implement a reverse proxy layer in front of web servers and configure IDS/ IPS agents on each reverse proxy server.

Answer: BD

NEW QUESTION 214
You are designing a photo sharing mobile app the application will store all pictures in a single Amazon 53 bucket.
Users will upload pictures from their mobile device directly to Amazon 53 and will be able to view and download their own pictures directly from Amazon 53.
You want to configure security to handle potentially millions of users in the most secure manner possible. What should your server-side application do when a new user registers on the photo sharing mobile application?

A. Create a set of long-term credentials using AWS Security Token Service with appropriate permissions Store these credentials in the mobile app and use them to access Amazon 53.
B. Record the user's Information in Amazon RDS and create a role in IAM with appropriate permission
C. When the user uses their mobile app create temporary credentials using the AWS Security Token Service 'Assume Role' function Store these credentials in the mobile app's memory and use them to access Amazon 53 Generate new credentials the next time the user runs the mobile app.
D. Record the user's Information in Amazon DynamoD
E. When the user uses their mobile app create temporary credentials using AWS Security Token Service with appropriate permissions Store these credentials in the mobile app's memory and use them to access Amazon 53 Generate new credentials the next time the user runs the mobile app.
F. Create IAM use
G. Assign appropriate permissions to the IAM user Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon 53.
H. Create an IAM use
I. Update the bucket policy with appropriate permissions for the IAM user Generate an access Key and secret Key for the IAM user, store them In the mobile app

and use these credentials to access Amazon 53.

Answer: B

NEW QUESTION 219
You have an application running on an EC2 Instance which will allow users to download fl ies from a private 53 bucket using a pre-assigned URL. Before generating the URL the application should verify the existence of the fi Ie in 53.
How should the application use AWS credentials to access the 53 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 53 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 53 bucke
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 53 bucke
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 222
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 tor 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 ad min the permission to manage the security groups and restrict other roles.

NEW QUESTION 225
You have a periodic Image analysis application that gets some files In Input analyzes them and tor each file writes some data in output to a ten file the number of files in input per day is high and concentrated in a few hours of the day.
Currently you have a server on EC2 with a large EBS volume that hosts the input data and the results it takes almost 20 hours per day to complete the process What services could be used to reduce the elaboration time and improve the availability of the solution?

A. 53 to store 1/0 file
B. SOS to distribute elaboration commands to a group of hosts working in paralle
C. Auto scaling to dynamically size the group of hosts depending on the length of the SOS queue
D. EBS with Provisioned IOPS (PIOPS) to store 1/0 file
E. SNS to distribute elaboration commands to a group of hosts working in parallel Auto Scaling to dynamically size the group of hosts depending on the number of SNS notifications
F. 53 to store 1/0 files, SNS to distribute evaporation commands to a group of hosts working in paralle
G. Auto scaling to dynamically size the group of hosts depending on the number of SNS notifications
H. EBS with Provisioned IOPS (PIOPS) to store 1/0 files SOS to distribute elaboration commands to a group of hosts working in parallel Auto Scaling to dynamically size the group ot hosts depending on the length of the SOS queue.

Answer: D

Explanation:
Amazon EBS allows you to create storage volumes and attach them to Amazon EC2 instances. Once attached, you can create a file system on top of these volumes, run a database, or use them in any other way you would use a block device. Amazon EBS volumes are placed in a specific Availability Zone, where they are automatically replicated to protect you from the failure of a single component.
Amazon EBS provides three volume types: General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic. The three volume types differ in performance characteristics and cost, so you can choose the right storage performance and price for the needs of your applications. All EBS volume types offer the same durable snapshot capabilities and are designed for 99.999% availability.

NEW QUESTION 229
A company is running a batch analysis every hour on their main transactional DB. running on an RDS MySQL instance to populate their central Data Warehouse running on Redshift During the execution of the batch their transactional applications are very slow When the batch completes they need to update the top management dashboard with the new data The dashboard is produced by another system running on-premises that is currently started when a manually-sent email notifies that an update is required The on-premises system cannot be modified because is managed by another team.
How would you optimize this scenario to solve performance issues and automate the process as much as possible?

A. Replace RDS with Redshift for the batch analysis and SNS to notify the on-premises system to update the dashboard
B. Replace ROS with Redshift for the oaten analysis and SQS to send a message to the on-premises system to update the dashboard
C. Create an RDS Read Replica for the batch analysis and SNS to notify me on-premises system to update the dashboard
D. Create an RDS Read Replica for the batch analysis and SQS to send a message to the on-premises system to update the dashboard

Answer: A

NEW QUESTION 234
You are implementing a URL whitelisting system for a company that wants to restrict outbound HTTP'S connections to specific domains from their EC2-hosted applications you deploy a single EC2 instance running proxy software and configure It to accept traffic from all subnets and EC2 instances in the VPC. You configure the proxy to only pass through traffic to domains that you define in its whitelist configuration You have a nightly maintenance window or 10 minutes where all instances fetch new software updates. Each update Is about 200MB In size and there are 500 instances In the VPC that routinely fetch updates After a few days you notice that some machines are failing to successfully download some, but not all of their updates within the maintenance window. The download URLs used for these updates are correctly listed in the proxy's whitelist configuration and you are able to access them manually using a web browser on the instances. What might be happening? {Choose 2 answers)

A. You are running the proxy on an undersized EC2 instance type so network throughput is not sufficient for all instances to download their updates in time.
B. You are running the proxy on a sufficiently-sized EC2 instance in a private subnet and its network throughput is being throttled by a NAT running on an undersized EC2 instance.
C. The route table for the subnets containing the affected EC2 instances is not configured to direct network traffic for the software update locations to the proxy.
D. You have not allocated enough storage to t he EC2 instance running the proxy so the network buffer is filling up, causing some requests to fail.
E. You are running the proxy in a public subnet but have not allocated enough EIPs to support the needed network throughput through the Internet Gateway {IGW).

Answer: AB

NEW QUESTION 238
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 variabl
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 241
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 accounts 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 Billin
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:Get0bject-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 t hat 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 al lows 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 tab Ie 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 244

You deployed your company website using Elastic Beanstalk and you enabled log file rotation to 53. An Elastic Map Reduce job is periodically analyzing the logs on 53 to build a usage dashboard that you share with your CIO.
You recently improved overall performance of the website using Cloud Front for dynamic content delivery and your website as the origin.
After this architectural change, the usage dashboard shows that the traffic on your website dropped by an order of magnitude. How do you fix your usage dashboard'?

A. Enable Cloud Front to deliver access logs to 53 and use them as input of the Elastic Map Reduce job.
B. Turn on Cloud Trail and use trail log tiles on 53 as input of the Elastic Map Reduce job
C. Change your log collection process to use Cloud Watch ELB metrics as input of the Elastic MapReducejob
D. Use Elastic Beanstalk "Rebuild Environment" option to update log delivery to the Elastic Map Reduce job.
E. Use Elastic Beanstalk 'Restart App server(s)" option to update log delivery to the Elastic Map Reduce job.

Answer: D

NEW QUESTION 245

A web-startup runs its very successful social news application on Amazon EC2 with an Elastic Load Balancer, an Auto-Scaling group of Java/Tomcat application-servers, and DynamoDB as data store. The main web-application best runs on m2 x large instances since it is highly memory- bound Each new deployment requires semi-automated creation and testing of a new AM for the application servers which takes quite a while ana is therefore only done once per week.
Recently, a new chat feature has been implemented in nodejs and wails to be integrated in the architecture. First tests show that the new component is CPU bound Because the company has some experience with using Chef, they decided to streamline the deployment process and use AWS Ops Works as an application life cycle tool to simplify management of the application and reduce the deployment cycles.
What configuration in AWS Ops Works is necessary to integrate the new chat module in the most cost-efficient and filexible way?

A. Create one AWS OpsWorks stack, create one AWS Ops Works layer, create one custom recipe
B. Create one AWS OpsWorks stack create two AWS Ops Works layers create one custom recipe
C. Create two AWS OpsWorks stacks create two AWS Ops Works layers create one custom recipe
D. Create two AWS OpsWorks stacks create two AWS Ops Works layers create two custom recipe

Answer: C

NEW QUESTION 249

What is the Reduced Redundancy option in Amazon 53?

A. Less redundancy for a lower cost.
B. It doesn't exist in Amazon 53, but in Amazon EBS.
C. It allows you to destroy any copy of your files outside a specific jurisdiction.
D. It doesn't exist at all

Answer: A

NEW QUESTION 251

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 253

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

A. No
B. Yes

Answer: B

NEW QUESTION 255

When you view the block device mapping for your instance, you can see only the EBS volumes, not the instance store volumes.

A. Depends on the instance type
B. FALSE
C. Depends on whether you use API call
D. TRUE

Answer: D

NEW QUESTION 256
True or False: When using IAM to control access to your RDS resources, the key names that can be used are case sensitive. For example, aws:CurrentTime is NOT equivalent to AWS:currenttime.

A. TRUE
B. FALSE

Answer: A

NEW QUESTION 259
Can we attach an EBS volume to more than one EC2 instance at the same time?

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

Answer: A

NEW QUESTION 263
True or False: Automated backups are enabled by default for a new DB Instance.

A. TRUE
B. FALSE

Answer: A

NEW QUESTION 267
Fill in the blanks: The base URI for all requests for instance metadata is _ _

A. http://254.169.169.254/latest/
B. http://169.169.254.254/|atesU
C. http://127.0.0.1/|atest/
D. http://I69.254.169.254/|atest/

Answer: D

NEW QUESTION 272
What are the two permission types used by AWS'?

A. Resource-based and Product-based
B. Product-based and Service-based
C. Service-based
D. User-based and Resource-based

Answer: D

NEW QUESTION 275
Changes to the backup window take effect _ _

A. from the next billing cycle
B. after 30 minutes
C. immediately
D. after 24 hours

Answer: C

NEW QUESTION 279
Is creating a Read Replica of another Read Replica supported?

A. Only in certain regions
B. Only with MSSQL based RDS
C. Only for Oracle RDS types
D. No

Answer: D

NEW QUESTION 281
What does the following command do with respect to the Amazon EC2 security groups? ec2-revoke RevokeSecurityGroup Ingress

A. Removes one or more security groups from a rule.
B. Removes one or more security groups from an Amazon EC2 instance.

C. Removes one or more rules from a security group.
D. Removes a security group from our accoun

Answer: C

NEW QUESTION 286
What does t he following command do with respect to the Amazon EC2 security groups? ec2-create-group CreateSecurityGroup

A. Groups the user created security groups in to a new group for easy access.
B. Creates a new security group for use with your account.
C. Creates a new group inside the security group.
D. Creates a new rule inside the security grou

Answer: B

NEW QUESTION 287
In the Launch Db Instance Wizard, where can I select the backup and maintenance options?

A. Under DB INSTANCE DETAILS
B. Under REVI EW
C. Under MANAGEMENT OPTIONS
D. Under ENGINE SELECTION

Answer: C

NEW QUESTION 292
For each DB Instance class, what is the maximum size of associated storage capacity?

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

Answer: B

NEW QUESTION 294
What happens to the 1/0 operations while you take a database snapshot?

A. 1/0 operations to the database are suspended for a few minutes while the backup is in progress.
B. 1/0 operations to the database are sent to a Replica (if available) for a few minutes while the backup is in progress.
C. 1/0 operations will be functioning normally
D. 1/0 operations to the database are suspended for an hour while the backup is in progress

Answer: A

NEW QUESTION 297
Which Amazon Storage behaves like raw, unformatted, external block devices that you can attach to your instances?

A. None of these.
B. Amazon Instance Storage
C. Amazon EBS
D. All of these

Answer: C

NEW QUESTION 298
Amazon RDS automated backups and DB Snapshots are currently supported for only the _ _ storage engine

A. MyISAM
B. InnoDB

Answer: B

NEW QUESTION 303
True or False: If you add a tag that has the same key as an existing tag on a DB Instance, the new value ovenrvrites the old value.

A. FALSE
B. TRUE

Answer: B

NEW QUESTION 306
What does Amazon Cloud Formation provide?

A. The ability to setup Autoscaling for Amazon EC2 instances.
B. None of these.
C. A templated resource creation for Amazon Web Services.
D. A template to map network resources for Amazon Web Service

Answer: D

NEW QUESTION 309
What does Amazon Route53 provide?

A. A global Content Delivery Network.
B. None of these.
C. A scalable Domain Name System.
D. An SSH endpoint for Amazon EC2.

Answer: C

NEW QUESTION 310
What does Amazon E|astiCache provide?

A. A senrice by this name doesn't exis
B. Perhaps you mean Amazon C|oudCache.
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-installe

Answer: C

NEW QUESTION 312
What is a Security Group?

A. None of these.
B. A list of users that can access Amazon EC2 instances.
C. An Access Control List (ACL) for AWS resources.
D. A firewall for inbound traffic, built-in around every Amazon EC2 instanc

Answer: D

NEW QUESTION 316
If an Amazon EBS volume is the root device of an instance, can I detach it without stopping the instance?

A. Yes but only if Windows instance
B. No
C. Yes
D. Yes but only if a Linux instance

Answer: B

NEW QUESTION 317
If you are using Amazon RDS Provisioned IOPS storage with MySQL and Oracle database engines, you can scale the throughput of your database Instance by specifying the IOPS rate f rom _ _

A. 1,000 to 1,00,000
B. 100 to 1, 000
C. 10,000to 1,00,000
D. 1, 000 to 10, 000

Answer: D

NEW QUESTION 319
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. Ye
C. You can.
D. N
E. You cannot.
F. Only if the tag "VPC Change Group" is true

Answer: B

NEW QUESTION 323
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

AWS-Solution-Architect-Associate Practice Exam Features:

* AWS-Solution-Architect-Associate Questions and Answers Updated Frequently

* AWS-Solution-Architect-Associate Practice Questions Verified by Expert Senior Certified Staff

* AWS-Solution-Architect-Associate Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Solution-Architect-Associate Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

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