SAA-C01 Dumps

AWS Certified Solutions Architect - Associate

https://www.certleader.com/SAA-C01-dumps.html

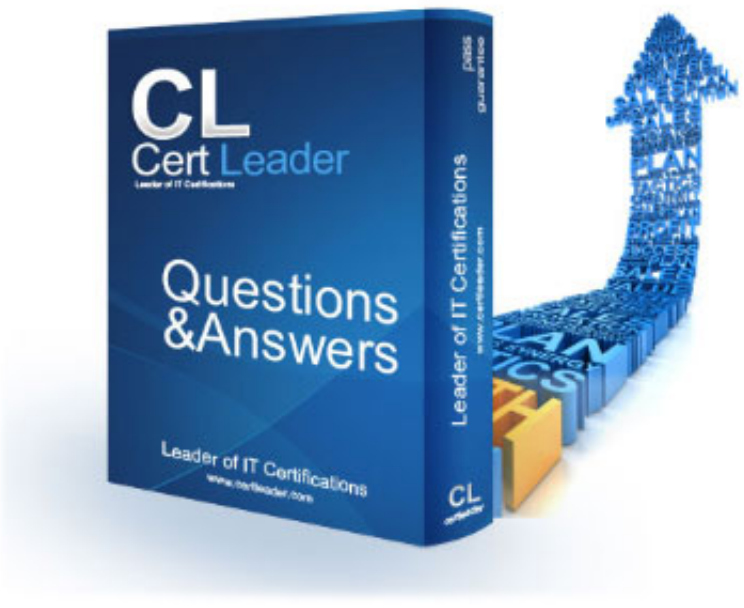

NEW QUESTION 1
A 3-tier e-commerce web application is current deployed on-premises and will be migrated to AWS for greater scalability and elasticity. The web server currently shares read-only data using a network distributed file system The app server tier uses a clustering mechanism for discovery and shared session state that depends on IP multicast The database tier uses shared-storage clustering to provide database fall over capability, and uses several read slaves for scaling Data on all servers and the distributed file system directory is backed up weekly to off-site tapes.
Which AWS storage and database architecture meets the requirements of the application?

A. Web servers: store read-only data in S3, and copy from S3 to root volume at boot tim
B. App servers: share state using a combination of DynamoDB and IP unicas
C. Database: use RDS with multi- AZ deployment and one or more read replica
D. Backup: web servers, app servers, and database backed up weekly to Glacier using snapshots.
E. Web servers: store read-only data in an EC2 NFS server, mount to each web server at boot tim
F. App servers: share state using a combination of DynamoDB and IP multicas
G. Database: use RDS with multi-AZ deployment and one or more Read Replica
H. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.
I. Web servers: store read-only data in S3, and copy from S3 to root volume at boot tim
J. App servers: share state using a combination of DynamoDB and IP unicas
K. Database: use RDS with multi- AZ deployment and one or more Read Replica
L. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.
M. Web servers: store read-only data in S3, and copy from S3 to root volume at boot tim
N. App servers: share state using a combination of DynamoDB and IP unicas
O. Database: use RDS with multi- AZ deploymen
P. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.

Answer: A

Explanation:
https://d0.awsstatic.com/whitepapers/Storage/AWS%20Storage%20Services%20Whitepaper-v9.pdf Amazon Glacier doesn't suit all storage situations. Listed following are a few storage needs for which you should consider other AWS storage options instead of Amazon Glacier.
Data that must be updated very frequently might be better served by a storage solution with lower read/write latencies, such as Amazon EBS, Amazon RDS, Amazon DynamoDB, or relational databases running on EC2.

NEW QUESTION 2
Your customer wishes to deploy an enterprise application to AWS which will consist of several web servers, several application servers and a small (50GB) Oracle database information is stored, both in the database and the file systems of the various servers. The backup system must support database recovery whole server and whole disk restores, and individual file restores with a recovery time of no more than two hours. They have chosen to use RDS Oracle as the database.
Which backup architecture will meet these requirements?

A. Backup RDS using automated daily DB backups Backup the EC2 instances using AMIs and supplement with file-level backup to S3 using traditional enterprise backup software to provide file level restore
B. Backup RDS using a Multi-AZ Deployment Backup the EC2 instances using Amis, and supplement by copying file system data to S3 to provide file level restore.
C. Backup RDS using automated daily DB backups Backup the EC2 instances using EBS snapshots and supplement with file-level backups to Amazon Glacier using traditional enterprise backup software to provide file level restore
D. Backup RDS database to S3 using Oracle RMAN Backup the EC2 instances using Amis, and supplement with EBS snapshots for individual volume restore.

Answer: A

Explanation:
You need to use enterprise backup software to provide file level restore. See
https://d0.awsstatic.com/whitepapers/Backup_and_Recovery_Approaches_Using_AWS.pdf Page 18:
If your existing backup software does not natively support the AWS cloud, you can use AWS storage gateway products. AWS Storage Gateway is a virtual appliance that provides seamless and secure integration between your data center and the AWS storage infrastructure.

NEW QUESTION 3
A customer has a 10 GB AWS Direct Connect connection to an AWS region where they have a web application hosted on Amazon Elastic Computer Cloud (EC2). The application has dependencies on an on-premises mainframe database that uses a BASE (Basic Available. Sort stale Eventual consistency) rather than an ACID (Atomicity. Consistency isolation. Durability) consistency model. The application is exhibiting undesirable behavior because the database is not able to handle the volume of writes. How can you reduce the load on your on-premises database resources in the most cost-effective way?

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

Answer: B

NEW QUESTION 4
Company B is launching a new game app for mobile devices. Users will log into the game using their existing social media account to streamline data capture. Company B would like to directly save player data and scoring information from the mobile app to a DynamoDS table named Score Data When a user saves their game the progress data will be stored to the Game state S3 bucket. What is the best approach for storing data to DynamoDB and S3?

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

with the mobile app.

Answer: B

Explanation:
The requirements state “Users will log into the game using their existing social media account to streamline data capture.” This is what Cognito is used for, ie Web Identity Federation. Amazon also recommend to “build your app so that it requests temporary AWS security credentials dynamically when needed using web identity federation.”

NEW QUESTION 5
You have recently joined a startup company building sensors to measure street noise and air quality in urban areas. The company has been running a pilot deployment of around 100 sensors for 3 months each sensor uploads 1KB of sensor data every minute to a backend hosted on AWS.
During the pilot, you measured a peak or 10 IOPS on the database, and you stored an average of 3GB of sensor data per month in the database.
The current deployment consists of a load-balanced auto scaled Ingestion layer using EC2 instances and a PostgreSQL RDS database with 500GB standard storage.
The pilot is considered a success and your CEO has managed to get the attention or some potential investors. The business plan requires a deployment of at least 100K sensors which needs to be
supported by the backend. You also need to store sensor data for at least two years to be able to compare year over year Improvements.
To secure funding, you have to make sure that the platform meets these requirements and leaves room for further scaling. Which setup win meet the requirements?

A. Add an SQS queue to the ingestion layer to buffer writes to the RDS instance
B. Ingest data into a DynamoDB table and move old data to a Redshift cluster
C. Replace the RDS instance with a 6 node Redshift cluster with 96TB of storage
D. Keep the current architecture but upgrade RDS storage to 3TB and 10K provisioned IOPS

Answer: C

Explanation:
You cannot go with DynamoDB because the application is currently using a Postgre SQL which is an RDS. Replacing an RDS SQL with a noSQL DB, for the sake of scaling is not a sensible option.

NEW QUESTION 6
You need a persistent and durable storage to trace call activity of an IVR (Interactive Voice Response)
system. Call duration is mostly in the 2-3 minutes timeframe. Each traced call can be either active or terminated. An external application needs to know each minute the list of currently active calls, which are usually a few calls/second. Put once per month there is a periodic peak up to 1000 calls/second for a few hours.
The system is open 24/7 and any downtime should be avoided. Historical data is periodically archived to files. Cost saving is a priority for this project.
What database implementation would better fit this scenario, keeping costs as low as possible?

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

Answer: B

Explanation:
Q: Can a global secondary index key be defined on non-unique attributes?
Yes. Unlike the primary key on a table, a GSI index does not require the indexed attributes to be unique.
Q: Are GSI key attributes required in all items of a DynamoDB table?
No. GSIs are sparse indexes. Unlike the requirement of having a primary key, an item in a DynamoDB table does not have to contain any of the GSI keys. If a GSI key has both hash and range elements, and a table item omits either of them, then that item will not be indexed by the corresponding GSI.
In such cases, a GSI can be very useful in efficiently locating items that have an uncommon attribute.

NEW QUESTION 7
A web design company currently runs several FTP servers that their 250 customers use to upload and download large graphic files. They wish to move this system to AWS to make it more scalable, but they wish to maintain customer privacy and Keep costs to a minimum.
What AWS architecture would you recommend?

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

Answer: A

Explanation:
In question we have keywords `scalable' and company wants to `move systems' to AWS, which is best suited for Auto-scaling group.
https://aws.amazon.com/blogs/security/writing-iam-policies-grant-access-to-user-specific-foldersin- an-amazon-s3-bucket/

NEW QUESTION 8
You would like to create a mirror image of your production environment in another region for disaster recovery purposes. Which of the following AWS resources do not need to be recreated in the second region? (Choose two.)

A. Route 53 Record Sets
B. IM1 Roles
C. Elastic IP Addresses (EIP)
D. EC2 Key Pairs
E. Launch configurations
F. Security Groups

Answer: AB

Explanation:
The Route 53 and IAM are global.
As per the document defined, new IPs should be reserved not the same ones. Elastic IP Addresses are static IP addresses designed for dynamic cloud computing. Unlike traditional static IP addresses, however, Elastic IP addresses enable you to mask instance or Availability Zone failures by programmatically remapping your public IP addresses to instances in your account in a particular region. For DR, you can also pre-allocate some IP addresses for the most critical systems so that their
IP addresses are already known before disaster strikes. This can simplify the execution of the DR plan.

NEW QUESTION 9
Your company currently has a 2-tier web application running in an on-premises data center. You have experienced several infrastructure failures in the past two months resulting in significant financial losses. Your CIO is strongly agreeing to move the application to AWS. While working on achieving buy-in from the other company executives, he asks you to develop a disaster recovery plan to help improve Business continuity in the short term. He specifies a target Recovery Time Objective (RTO) of 4 hours and a Recovery Point Objective (RPO) of 1 hour or less. He also asks you to implement the solution within 2 weeks. Your database is 200GB in size and you have a 20Mbps Internet connection. How would you do this while minimizing costs?

A. Create an EBS backed private AMI which includes a fresh install of your applicatio
B. Develop a CloudFormation template which includes your AMI and the required EC2, AutoScaling, and ELB resources to support deploying the application across Multiple- Availability-Zone
C. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
D. Deploy your application on EC2 instances within an Auto Scaling group across multiple availability zone
E. Asynchronously replicate transactions from your on-premises database to a database instancein AWS across a secure VPN connection.
F. Create an EBS backed private AMI which includes a fresh install of your applicatio
G. Setup a script in your data center to backup the local database every 1 hour and to encrypt and copy the resultingfile to an S3 bucket using multi-part upload.
H. Install your application on a compute-optimized EC2 instance capable of supporting the application's average loa
I. Synchronously replicate transactions from your on-premises database to a database instance in AWS across a secure Direct Connect connection.

Answer: A

Explanation:
Overview of Creating Amazon EBS-Backed AMIs
First, launch an instance from an AMI that's similar to the AMI that you'd like to create. You can
connect to your instance and customize it. When the instance is configured correctly, ensure data integrity by stopping the instance before you create an AMI, then create the image. When you create an Amazon EBS-backed AMI, we automatically register it for you.
Amazon EC2 powers down the instance before creating the AMI to ensure that everything on the instance is stopped and in a consistent state during the creation process. If you're confident that your instance is in a consistent state appropriate for AMI creation, you can tell Amazon EC2 not to power down and reboot the instance. Some file systems, such as XFS, can freeze and unfreeze activity, making it safe to create the image without rebooting the instance.
During the AMI-creation process, Amazon EC2 creates snapshots of your instance's root volume and any other EBS volumes attached to your instance. If any volumes attached to the instance are encrypted, the new AMI only launches successfully on instances that support Amazon EBS encryption. For more information, see Amazon EBS Encryption.
Depending on the size of the volumes, it can take several minutes for the AMI-creation process to complete (sometimes up to 24 hours). You may find it more efficient to create snapshots of your volumes prior to creating your AMI. This way, only small, incremental snapshots need to be created when the AMI is created, and the process completes more quickly (the total time for snapshot creation remains the same). For more information, see Creating an Amazon EBS Snapshot.
After the process completes, you have a new AMI and snapshot created from the root volume of the instance. When you launch an instance using the new AMI, we create a new EBS volume for its root volume using the snapshot. Both the AMI and the snapshot incur charges to your account until you delete them. For more information, see Deregistering Your AMI.
If you add instance-store volumes or EBS volumes to your instance in addition to the root device volume, the block device mapping for the new AMI contains information for these volumes, and the block device mappings for instances that you launch from the new AMI automatically contain information for these volumes. The instance-store volumes specified in the block device mapping for the new instance are new and don't contain any data from the instance store volumes of the instance you used to create the AMI. The data on EBS volumes persists. For more information, see Block Device Mapping.

NEW QUESTION 10
You have deployed a web application targeting a global audience across multiple AWS Regions under the domain name.example.com. You decide to use Route53 Latency-Based Routing to serve web requests to users from the region closest to the user. To provide business continuity in the event of server downtime you configure weighted record sets associated with two web servers in separate Availability Zones per region. Running a DR test you notice that when you disable all web servers in one of the regions Route53 does not automatically direct all users to the other region. What could be happening? (Choose two.)

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

Answer: BE

Explanation:
How Health Checks Work in Complex Amazon Route 53 Configurations

Checking the health of resources in complex configurations works much the same way as in simple configurations. However, in complex configurations, you use a combination of alias resource record sets (including weighted alias, latency alias, and failover alias) and nonalias resource record sets to build a decision tree that gives you greater control over how Amazon Route 53 responds to requests. For more information, see How Health Checks Work in Simple Amazon Route 53 Configurations. For example, you might use latency alias resource record sets to select a region close to a user and use weighted resource record sets for two or more resources within each region to protect against the failure of a single endpoint or an Availability Zone. The following diagram shows this configuration.

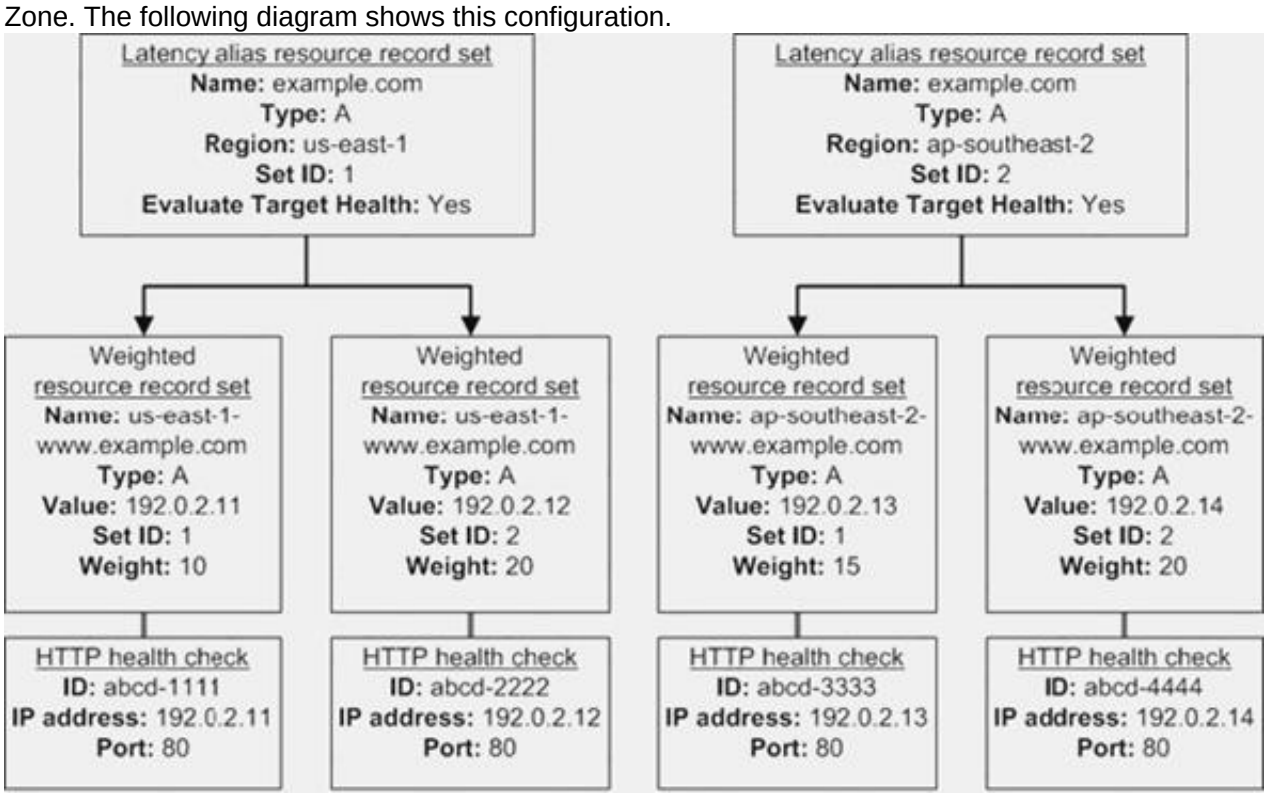

Here's how Amazon EC2 and Amazon Route 53 are configured:

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

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

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

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

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

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

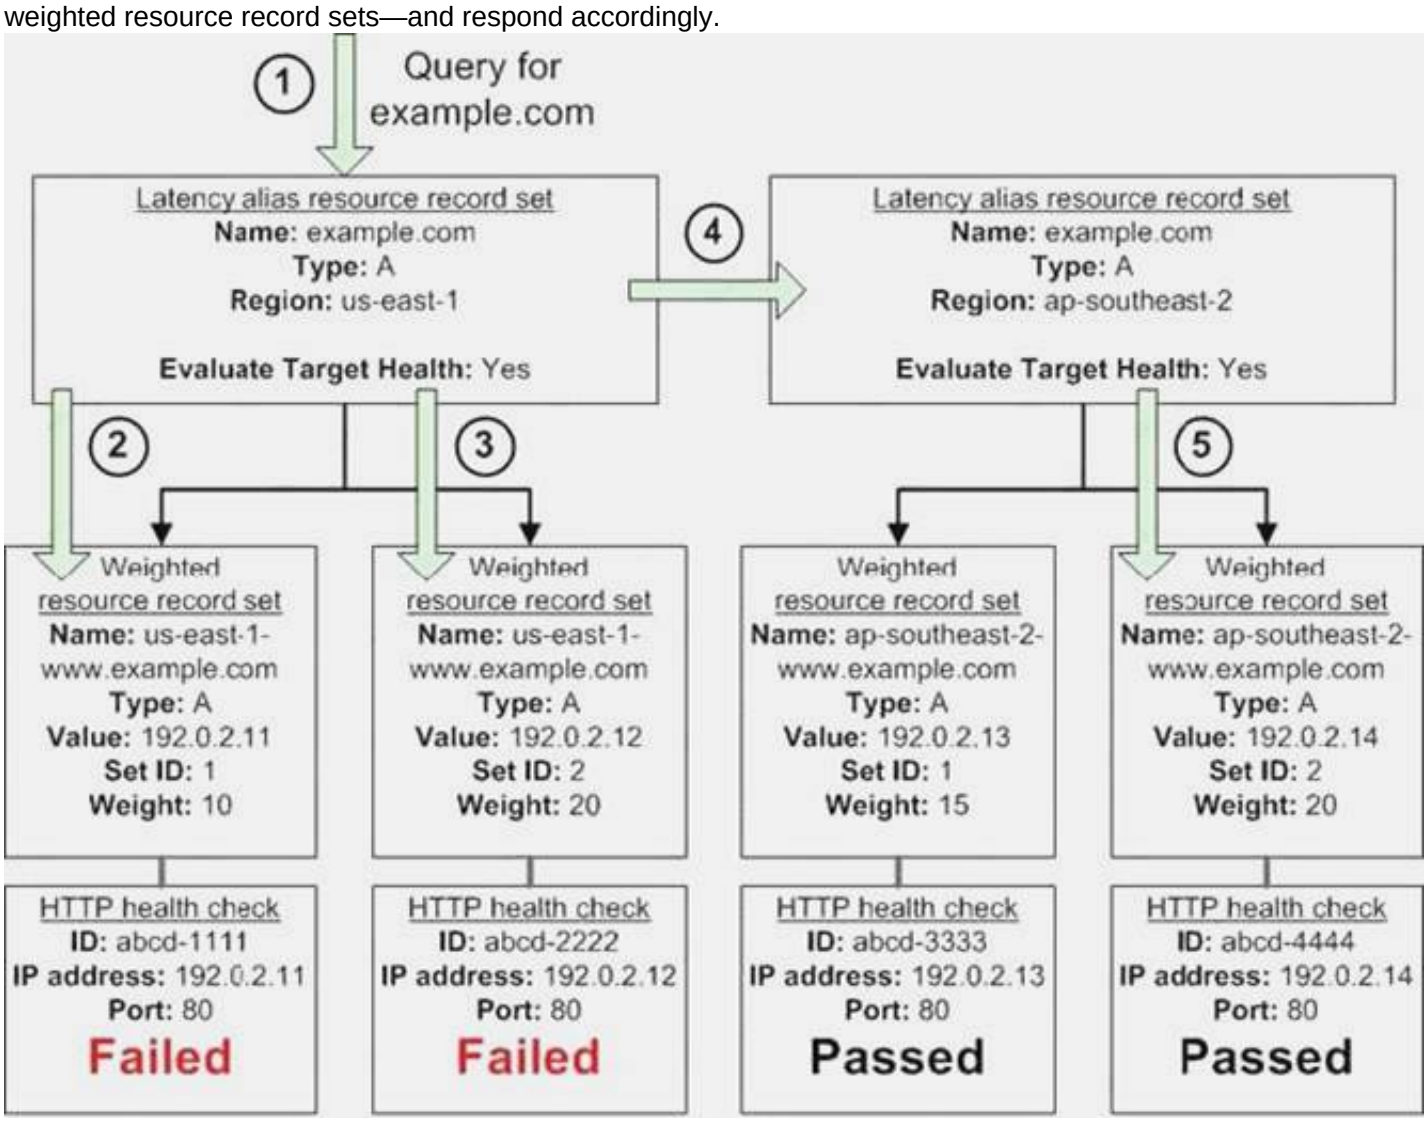

The preceding diagram illustrates the following sequence of events:

Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region. Amazon Route 53 selects a weighted resource record set based on weight. Evaluate Target Health is Yes for the latency alias resource record set, so Amazon Route 53 checks the health of the selected weighted resource record set.

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

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

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

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

You can associate a health check with an alias resource record set instead of or in addition to setting the value of Evaluate Target Health to Yes. However, it's generally more useful if Amazon Route 53 responds to queries based on the health of the underlying resources—the HTTP servers, database servers, and other resources that your alias resource record sets refer to. For example, suppose the following configuration:
You assign a health check to a latency alias resource record set for which the alias target is a group of weighted resource record sets.
You set the value of Evaluate Target Health to Yes for the latency alias resource record set.
In this configuration, both of the following must be true before Amazon Route 53 will return the applicable value for a weighted resource record set:
The health check associated with the latency alias resource record set must pass.
At least one weighted resource record set must be considered healthy, either because it's associated with a health check that passes or because it's not associated with a health check. In the latter case, Amazon Route 53 always considers the weighted resource record set healthy.

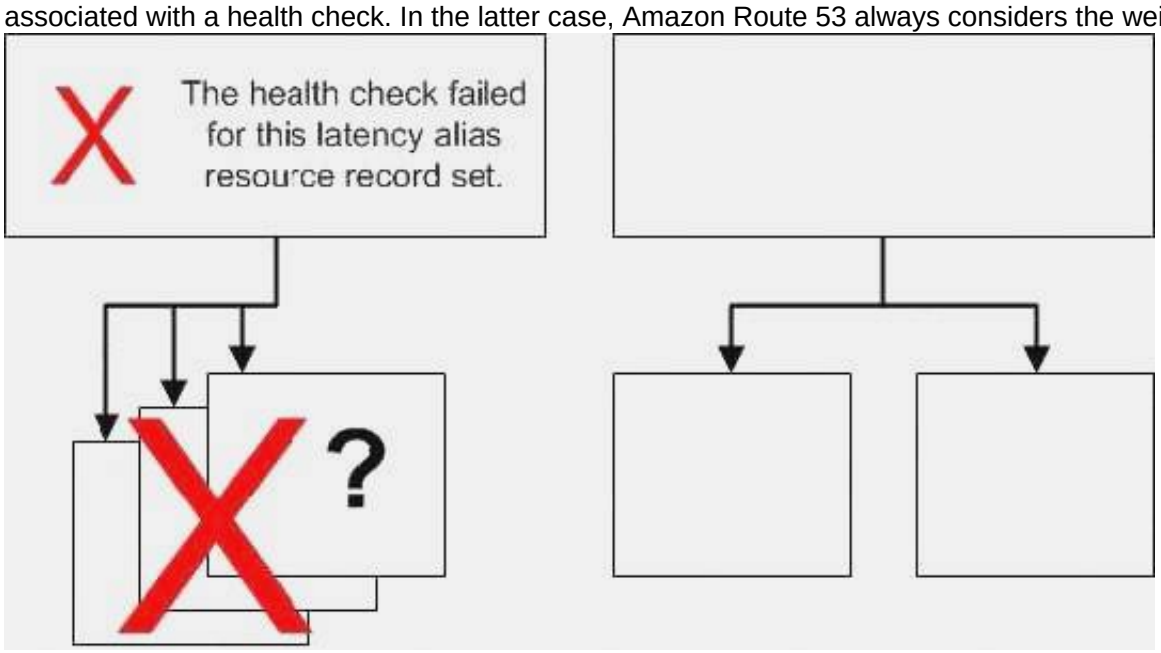

If the health check for the latency alias resource record set fails, Amazon Route 53 stops responding to queries using any of the weighted resource record sets in the alias target, even if they're all healthy. Amazon Route 53 doesn't know the status of the weighted resource record sets because it never looks past the failed health check on the alias resource record set.
What Happens When You Omit Health Checks?
In a complex configuration, it's important to associate health checks with all of the non-alias resource record sets. Let's return to the preceding example, but assume that a health check is missing on one of the weighted resource record sets in the us-east-1 region:

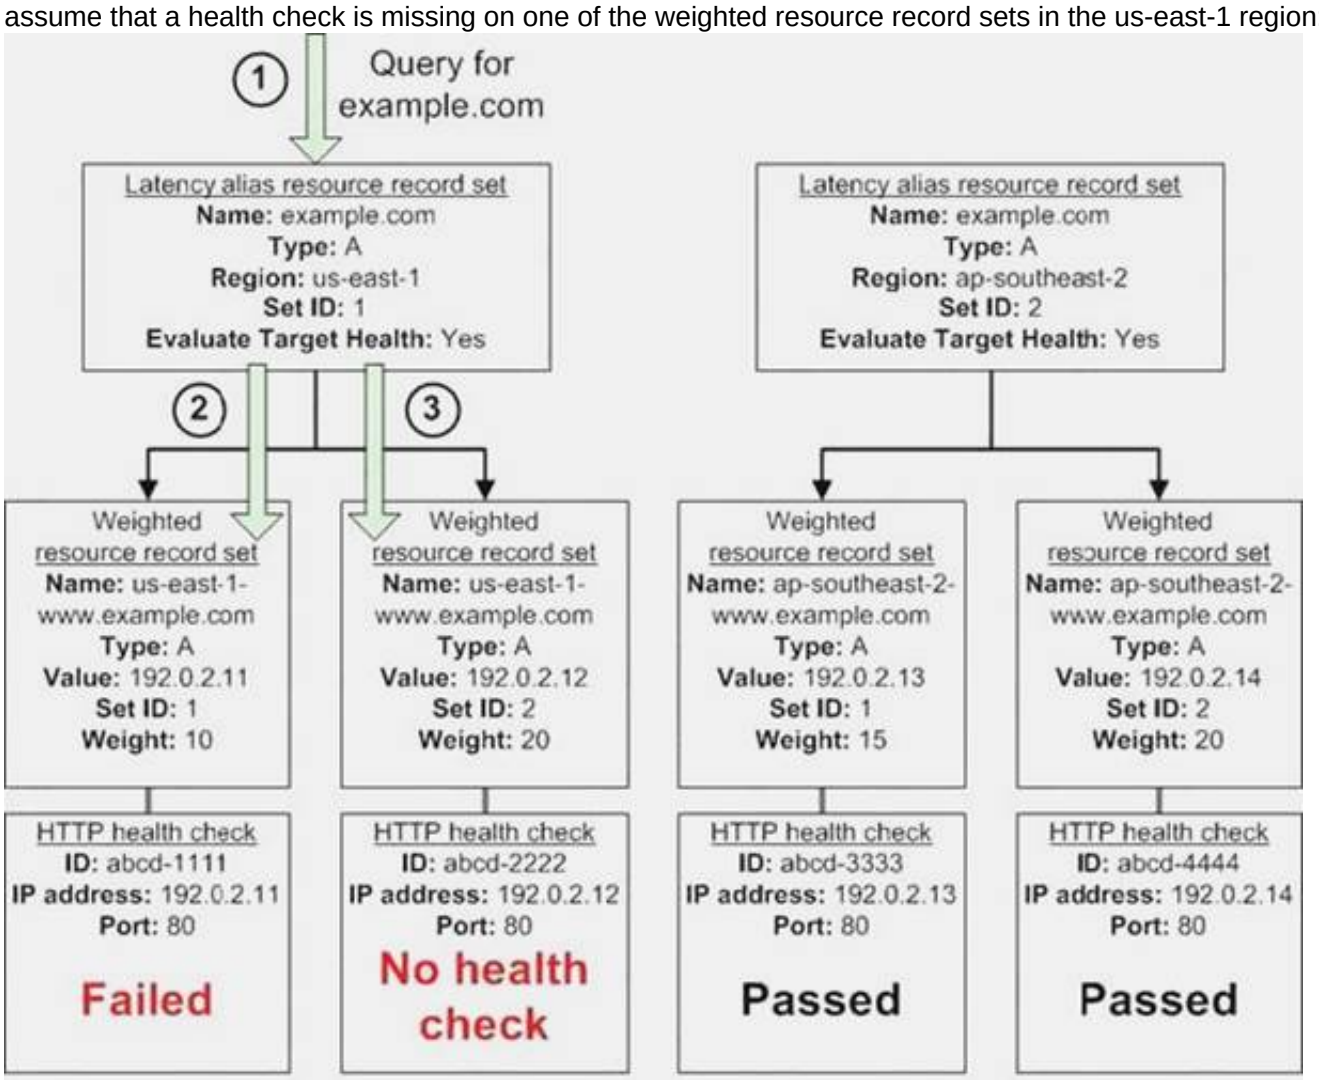

Here's what happens when you omit a health check on a non-alias resource record set in this configuration:
Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region. Amazon Route 53 looks up the alias target for the latency alias resource record set, and checks the status of the corresponding health checks. The health check for one weighted resource record set failed, so that resource record set is omitted from consideration.
The other weighted resource record set in the alias target for the us-east-1 region has no health check. The corresponding resource might or might not be healthy, but without a health check, Amazon Route 53 has no way to know. Amazon Route 53 assumes that the resource is healthy and returns the applicable value in response to the query.
What Happens When You Set Evaluate Target Health to No?
In general, you also want to set Evaluate Target Health to Yes for all of the alias resource record sets. In the following example, all of the weighted resource record sets have associated health checks, but Evaluate Target Health is set to No for the latency alias resource record set for the us-east-1 region:

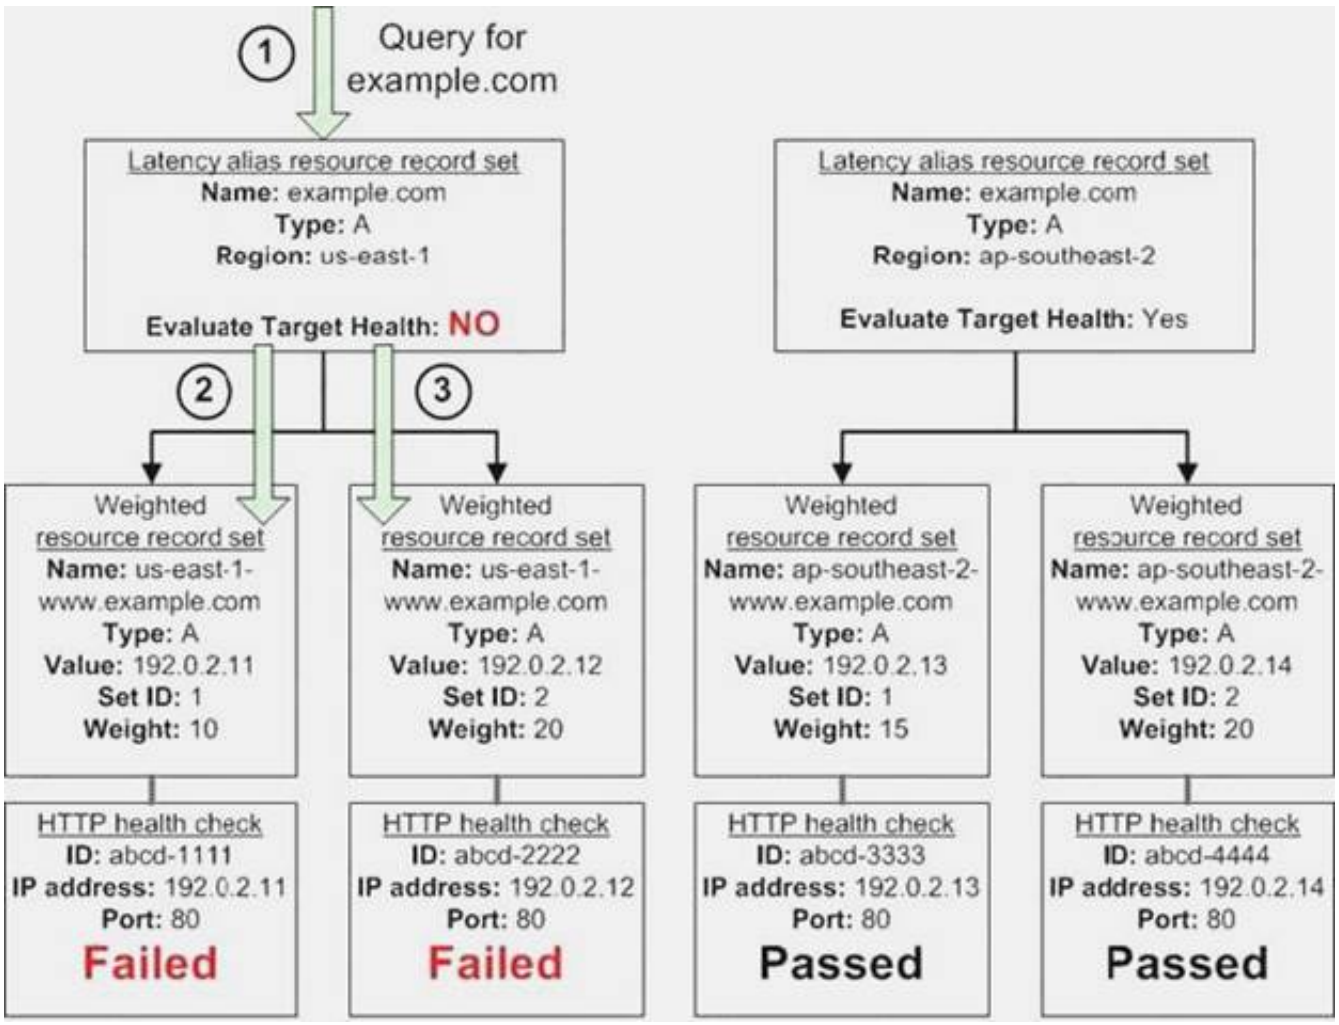

Here's what happens when you set Evaluate Target Health to No for an alias resource record set in this configuration:
Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region. Amazon Route 53 determines what the alias target is for the latency alias resource record set, and checks the corresponding health checks. They're both failing.
Because the value of Evaluate Target Health is No for the latency alias resource record set for the useast- 1 region, Amazon Route 53 must choose one resource record set in this branch instead of
backing out of the branch and looking for a healthy resource record set in the ap-southeast-2 region.

NEW QUESTION 10
Your company hosts a social media site supporting users in multiple countries. You have been asked to provide a highly available design tor the application that leverages multiple regions tor the most recently accessed content and latency sensitive portions of the wet) site The most latency sensitive component of the application involves reading user preferences to support web site personalization and ad selection.
In addition to running your application in multiple regions, which option will support this
application's requirements?

A. Serve user content from S3. CloudFront and use Route53 latency-based routing between ELBs in each region Retrieve user preferences from a local DynamoDB table in each region and leverage SQS to capture changes to user preferences with SOS workers for propagating updates to each table.
B. Use the S3 Copy API to copy recently accessed content to multiple regions and serve user content from S3. CloudFront with dynamic content and an ELB in each region Retrieve user preferences from an ElasticCache cluster in each region and leverage SNS notifications to propagate user preference changes to a worker node in each region.
C. Use the S3 Copy API to copy recently accessed content to multiple regions and serve user content from S3 CloudFront and Route53 latency-based routing Between ELBs In each region Retrieve user preferences from a DynamoDB table and leverage SQS to capture changes to user preferences with SOS workers for propagating DynamoDB updates.
D. Serve user content from S3. CloudFront with dynamic content, and an ELB in each region Retrieve user preferences from an ElastiCache cluster in each region and leverage Simple Workflow (SWF) to manage the propagation of user preferences from a centralized OB to each ElastiCache cluster.

Answer: A

Explanation:
http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_mediasharing_09.pdf
http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_adserving_06.pdf

NEW QUESTION 13
A customer has established an AWS Direct Connect connection to AWS. The link is up and routes are being advertised from the customer's end, however the customer is unable to connect from EC2 instances inside its VPC to servers residing in its datacenter.
Which of the following options provide a viable solution to remedy this situation? (Choose two.)

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

Answer: BE

NEW QUESTION 16
A web company is looking to implement an external payment service into their highly available application deployed in a VPC. Their application EC2 instances are behind a public lacing ELB Auto scaling is used to add additional instances as traffic increases under normal load the application runs 2 instances in the Auto Scaling group but at peak it can scale 3x in size. The application instances need to communicate with the payment service over the Internet which requires whitelisting of all public IP addresses used to communicate with it. A maximum of 4 whitelisting IP addresses is allowed at a time and can be added through an

API.
How should they architect their solution?

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

Answer: A

Explanation:
B is incorrect as you do not have insight into the public ip associated with a VPC Internet Gateways. C is incorrect as ELB receives a public DNS name.
D would exceed the maximum of 4 whitelisting IP addresses.

NEW QUESTION 18
You are designing the network infrastructure for an application server in Amazon VPC Users will
access all the application instances from the Internet as well as from an on-premises network The onpremises network is connected to your VPC over an AWS Direct Connect link.
How would you design routing to meet the above requirements?

A. Configure a single routing Table with a default route via the Internet gateway Propagate a default route via BGP on the AWS Direct Connect customer route
B. Associate the routing table with all VPC subnets.
C. Configure a single routing table with a default route via the internet gateway Propagate specific routes for the on-premises networks via BGP on the AWS Direct Connect customer router Associate the routing table with all VPC subnets.
D. Configure a single routing table with two default routes: one to the internet via an Internet gateway the other to the on-premises network via the VPN gateway use this routing table across all subnets in your VPC,
E. Configure two routing tables one that has a default route via the Internet gateway and another that has a default route via the VPN gateway Associate both routing tables with each VPC subnet.

Answer: B

NEW QUESTION 21
You have deployed a three-tier web application in a VPC with a CIOR block of 10 0 0 0/28. You initially deploy two web servers, two application servers, two database servers and one NAT instance tor a
total of seven EC2 instances. The web. Application and database servers are deployed across two availability zones (AZs). You also deploy an ELB in front of the two web servers, and use Route53 for DNS Web (raffile gradually increases in the first few days following the deployment, so you attempt to double the number of instances in each tier of the application to handle the new load unfortunately some of these new instances fail to launch.
Which of the following could De the root caused? (Choose two.)

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

Answer: CE

Explanation:
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

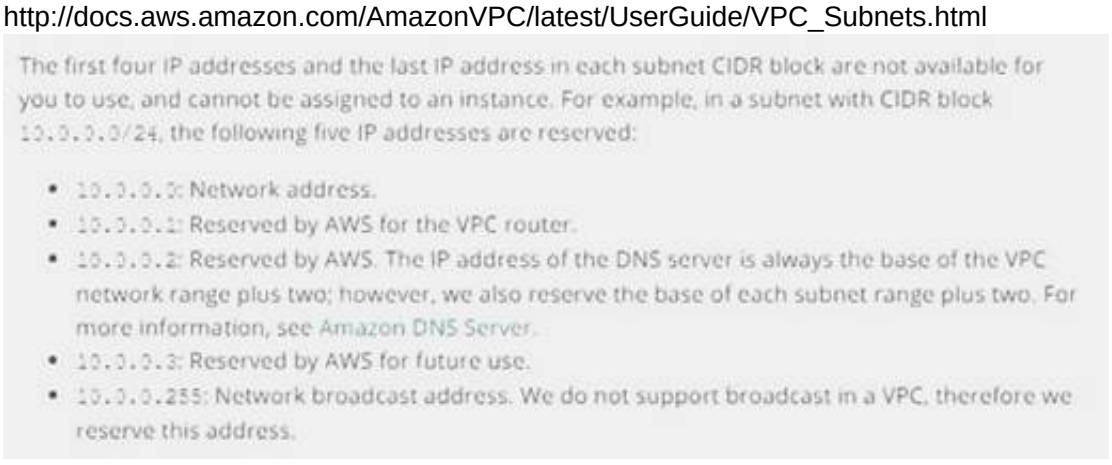
The first four IP addresses and the last IP address in each subnet CIDR block are not available for you to use, and cannot be assigned to an instance. For example, in a subnet with CIDR block 10.0.0.0/24, the following five IP addresses are reserved:

- 10.0.0.0: Network address.
- 10.0.0.1: Reserved by AWS for the VPC router.
- 10.0.0.2: Reserved by AWS. The IP address of the DNS server is always the base of the VPC network range plus two; however, we also reserve the base of each subnet range plus two. For more information, see Amazon DNS Server.
- 10.0.0.3: Reserved by AWS for future use.
- 10.0.0.255: Network broadcast address. We do not support broadcast in a VPC, therefore we reserve this address.

NEW QUESTION 26
You've been brought in as solutions architect to assist an enterprise customer with their migration of an e-commerce platform to Amazon Virtual Private Cloud (VPC) The previous architect has already
deployed a 3-tier VPC,
The configuration is as follows:

```
VPC: vpc-2f8bc447
IGW: igw-2d8bc445
NACL: ad-208bc448
Subnets and Route Tables:
Web servers: subnet-258bc44d
Application servers: subnet-248bc44c
Database servers: subnet-9189c6f9
Route Tables:
rrb-218bc449
rtb-238bc44b
Associations:
subnet-258bc44d : rtb-218bc449
subnet-248bc44c : rtb-238bc44b
subnet-9189c6f9 : rtb-238bc44b
```

You are now ready to begin deploying EC2 instances into the VPC. Web servers must have direct access to the internet. Application and database servers cannot have direct access to the internet. Which configuration below will allow you the ability to remotely administer your application and database servers, as well as allow these servers to retrieve updates from the Internet?

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

Answer: A

Explanation:
Create NAT instance in public subnet which is web server subnet (suDnet-258Dc44d) and add route (rtD-238Dc44D) from private subnet (database subnet-9189c6f9) to the public NAT one to
retrieve the updates.

NEW QUESTION 27
You are designing Internet connectivity for your VPC. The Web servers must be available on the Internet. The application must have a highly available architecture.
Which alternatives should you consider? (Choose two.)

A. Configure a NAT instance in your VP
B. Create a default route via the NAT instance and associate it with all subnet
C. Configure a DNS A record that points to the NAT instance public IP address.
D. Configure a CloudFront distribution and configure the origin to point to the private IP addresses ofyour Web server
E. Configure a Route53 CNAME record to your CloudFront distribution.
F. Place all your web servers behind EL8. Configure a Route53 CNMIE to point to the ELB DNS name.
G. Assign EIPs to all web server
H. Configure a Route53 record set with all EIPs with health checks and DNS failover.
I. Configure ELB with an EI
J. Place all your Web servers behind EL
K. Configure a Route53 A record that points to the EIP.

Answer: CD

NEW QUESTION 32
You are tasked with moving a legacy application from a virtual machine running Inside your datacenter to an Amazon VPC. Unfortunately this app requires access to a number of on-premises services and no one who configured the app still works for your company. Even worse there's no documentation for it. What will allow the application running inside the VPC to reach back and access its internal dependencies without being reconfigured? (Choose three.)

A. An AWS Direct Connect link between the VPC and the network housing the internal services.
B. An Internet Gateway to allow a VPN connection.
C. An Elastic IP address on the VPC instance
D. An IP address space that does not conflict with the one on-premises
E. Entries in Amazon Route 53 that allow the Instance to resolve its dependencies' IP addresses
F. A VM Import of the current virtual machine

Answer: ADF

Explanation:
AWS Direct Connect
AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your datacenter, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internetbased connections.
AWS Direct Connect lets you establish a dedicated network connection between your network and one of the AWS Direct Connect locations. Using industry standard 802.1q VLANs, this dedicated connection can be partitioned into multiple virtual interfaces. This allows you to use the same connection to access public resources such as objects stored in Amazon S3 using public IP address space, and private resources such as Amazon EC2 instances running within an Amazon Virtual Private Cloud (VPC) using private IP space, while maintaining network separation between the public and private environments. Virtual interfaces can be reconfigured at any time to meet your changing
needs.
What is AWS Direct Connect?
AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 1 gigabit or 10 gigabit Ethernet fiber-optic cable. One end of the cable is connected to your router, the other to an AWS Direct Connect router. With this connection in place, you can create virtual interfaces directly to the AWS cloud (for example, to Amazon Elastic Compute Cloud (Amazon EC2) and Amazon Simple Storage Service (Amazon S3)) and to Amazon Virtual Private Cloud (Amazon VPC), bypassing Internet service providers in your network path. An AWS Direct Connect location provides access to Amazon Web Services in the

region it is associated with, as well as access to other
US regions. For example, you can provision a single connection to any AWS Direct Connect location in the US and use it to access public AWS services in all US Regions and AWS GovCloud (US).
The following diagram shows how AWS Direct Connect interfaces with your network.

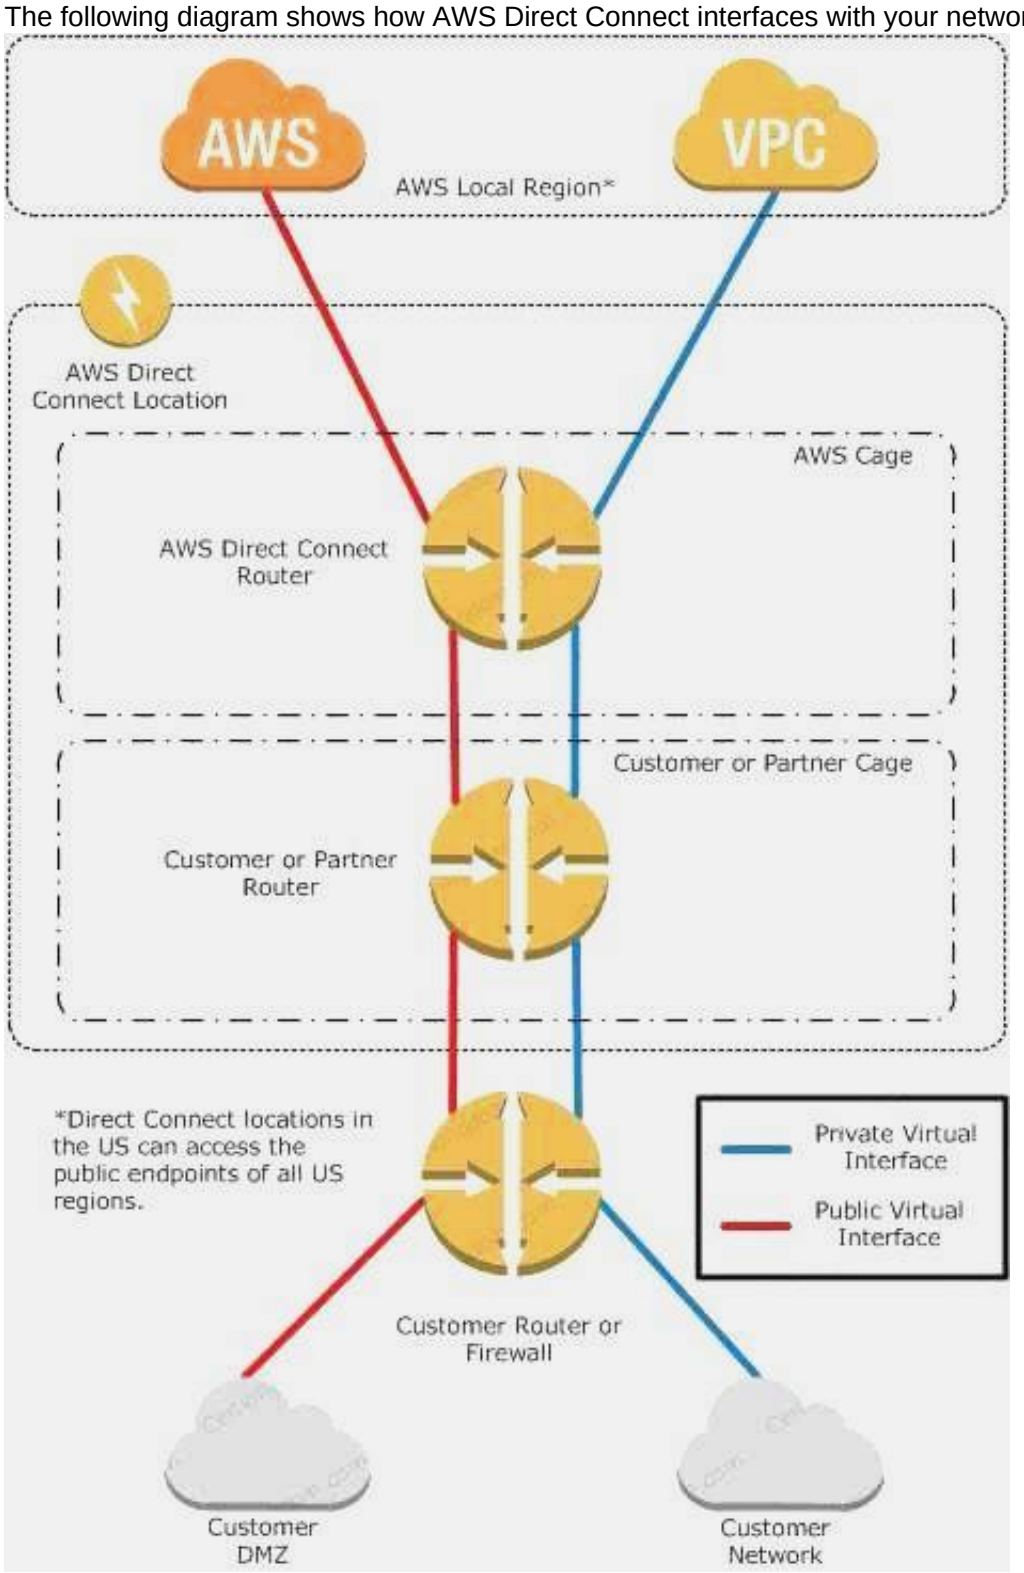

Requirements
To use AWS Direct Connect, your network must meet one of the following conditions:
Your network is colocated with an existing AWS Direct Connect location. For more information on
available AWS Direct Connect locations, go to http://aws.amazon.com/directconnect/.
You are working with an AWS Direct Connect partner who is a member of the AWS Partner Network (APN). For a list of AWS Direct Connect partners who can help you connect, go to http://aws.amazon.com/directconnect.
You are working with an independent service provider to connect to AWS Direct Connect. In addition, your network must meet the following conditions:
Connections to AWS Direct Connect require single mode fiber, 1000BASE-LX (1310nm) for 1 gigabit Ethernet, or 10GBASE-LR (1310nm) for 10 gigabit Ethernet.
Auto Negotiation for the port must be disabled. You must support 802.1Q VLANs across these connections.
Your network must support Border Gateway Protocol (BGP) and BGP MD5 authentication. Optionally, you may configure Bidirectional Forwarding Detection (BFD).
To connect to Amazon Virtual Private Cloud (Amazon VPC), you must first do the following: Provide a private Autonomous System Number (ASN). Amazon allocates a private IP address in the
169.x.x.x range to you.
Create a virtual private gateway and attach it to your VPC. For more information about creating a virtual private gateway, see Adding a Hardware Virtual Private Gateway to Your VPC in the Amazon VPC User Guide.
To connect to public AWS products such as Amazon EC2 and Amazon S3, you need to provide the following:
A public ASN that you own (preferred) or a private ASN.
Public IP addresses (/31) (that is, one for each end of the BGP session) for each BGP session. If you do not have public IP addresses to assign to this connection, log on to AWS and then open a ticket with AWS Support.
The public routes that you will advertise over BGP.

NEW QUESTION 36
A newspaper organization has an on-premises application, which allows the public to search its back catalogue and retrieve individual newspaper pages via a website written in Java They have scanned the old newspapers into JPEGs (approx 17TB) and used Optical Character Recognition (OCR) to populate a commercial search product. The hosting platform and software are now end of life and the organization wants to migrate Its archive to AWS and produce a cost efficient architecture and still be designed for availability and durability. Which is the most appropriate?

A. Use S3 with reduced redundancy lo store and serve the scanned files, install the commercial search application on EC2 Instances and configure with auto-scaling and an Elastic Load Balancer.
B. Model the environment using CloudFormation use an EC2 instance running Apache webserver and an open source search application, stripe multiple standard EBS volumes together to store the JPEGs and search index.
C. Use S3 with standard redundancy to store and serve the scanned files, use CloudSearch for query processing, and use Elastic Beanstalk to host the website

across multiple availability zones.
D. Use a single-AZ RDS MySQL instance lo store the search index 33d the JPEG images use an EC2 instance to serve the website and translate user queries into SQL.
E. Use a CloudFront download distribution to serve the JPEGs to the end users and Install the current commercial search product, along with a Java Container Tor the website on EC2 instances and use Route53 with DNS round-robin.

Answer: C

Explanation:
There is no such thing as "Most appropriate" without knowing all your goals. I find your scenarios very fuzzy, since you can obviously mix-n-match between them. I think you should decide by layers instead:
Load Balancer Layer: ELB or just DNS, or roll-your-own. (Using DNS+EIPs is slightly cheaper, but less reliable than ELB.)
Storage Layer for 17TB of Images: This is the perfect use case for S3. Off-load all the web requests directly to the relevant JPEGs in S3. Your EC2 boxes just generate links to them.
If your app already serves it's own images (not links to images), you might start with EFS. But more than likely, you can just setup a web server to re-write or re-direct all JPEG links to S3 pretty easily. If you use S3, don't serve directly from the bucket - Serve via a CNAME in domain you control. That way, you can switch in CloudFront easily.
EBS will be way more expensive, and you'll need 2x the drives if you need 2 boxes. Yuck. Consider a smaller storage format. For example, JPEG200 or WebP or other tools might make for smaller images. There is also the DejaVu format from a while back.
Cache Layer: Adding CloudFront in front of S3 will help people on the other side of the world -- well, possibly. Typical archives follow a power law. The long tail of requests means that most JPEGs won't be requested enough to be in the cache. So you are only speeding up the most popular objects. You can always wait, and switch in CF later after you know your costs better. (In some cases, it can actually lower costs.)
You can also put CloudFront in front of your app, since your archive search results should be fairly static. This will also allow you to run with a smaller instance type, since CF will handle much of the
load if you do it right. Database Layer: A few options:
Use whatever your current server does for now, and replace with something else down the road. Don't under-estimate this approach, sometimes it's better to start now and optimize later.
Use RDS to run MySQL/Postgres
I'm not as familiar with ElasticSearch / Cloudsearch, but obviously Cloudsearch will be less maintenance+setup.
App Layer:
When creating the app layer from scratch, consider CloudFormation and/or OpsWorks. It's extra stuff to learn, but helps down the road.
Java+Tomcat is right up the alley of ElasticBeanstalk. (Basically EC2 + Autoscale + ELB). Preventing Abuse: When you put something in a public S3 bucket, people will hot-link it from their web pages. If you want to prevent that, your app on the EC2 box can generate signed links to S3 that expire in a few hours. Now everyone will be forced to go thru the app, and the app can apply rate limiting, etc.
Saving money: If you don't mind having downtime:
run everything in one AZ (both DBs and EC2s). You can always add servers and AZs down the road, as long as it's architected to be stateless. In fact, you should use multiple regions if you want it to be
really robust.
use Reduced Redundancy in S3 to save a few hundred bucks per month (Someone will have to "go fix it" every time it breaks, including having an off-line copy to repair S3.)
Buy Reserved Instances on your EC2 boxes to make them cheaper. (Start with the RI market and buy a partially used one to get started.) It's just a coupon saying "if you run this type of box in this AZ, you will save on the per-hour costs." You can get 1/2 to 1/3 off easily.
Rewrite the application to use less memory and CPU - that way you can run on fewer/smaller boxes. (May or may not be worth the investment.)
If your app will be used very infrequently, you will save a lot of money by using Lambda. I'd be worried that it would be quite slow if you tried to run a Java application on it though.
We're missing some information like load, latency expectations from search, indexing speed, size of the search index, etc. But with what you've given us, I would go with S3 as the storage for the files (S3 rocks. It is really, really awesome). If you're stuck with the commercial search application, then on EC2 instances with autoscaling and an ELB. If you are allowed an alternative search engine, Elasticsearch is probably your best bet. I'd run it on EC2 instead of the AWS Elasticsearch service, as IMHO it's not ready yet. Don't autoscale Elasticsearch automatically though, it'll cause all sorts of issues. I have zero experience with CloudSearch so ic an't comment on that. Regardless of which option, I'd use CloudFormation for all of it.

NEW QUESTION 37
A corporate web application is deployed within an Amazon Virtual Private Cloud (VPC) and is connected to the corporate data center via an iPsec VPN. The application must authenticate against the on-premises LDAP server. After authentication, each logged-in user can only access an Amazon Simple Storage Space (S3) keyspace specific to that user.
Which two approaches can satisfy these objectives? (Choose two.)

A. Develop an identity broker that authenticates against IAM security Token service to assume a IAM role in order to get temporary AWS security credentials The application calls the identity broker to get AWS temporary security credentials with access to the appropriate S3 bucket.
B. The application authenticates against LDAP and retrieves the name of an IAM role associated withthe use
C. The application then calls the IAM Security Token Service to assume that IAM rol
D. The application can use the temporary credentials to access the appropriate S3 bucket.
E. Develop an identity broker that authenticates against LDAP and then calls IAM Security Token Service to get IAM federated user credential
F. The application calls the identity broker to get IAM federated user credentials with access to the appropriate S3 bucket.
G. The application authenticates against LDAP the application, then calls the AWS identity and Access Management (IAM) Security service to log in to IAM using the LDAP credentials, the application can use the IAM temporary credentials to access the appropriate S3 bucket.
H. The application authenticates against IAM Security Token Service using the LDAP credentials, the application uses those temporary AWS security credentials to access the appropriate S3 bucket.

Answer: BC

Explanation:
Imagine that in your organization, you want to provide a way for users to copy data from their computers to a backup folder. You build an application that users can run on their computers. On the back end, the application reads and writes objects in an S3 bucket. Users don’t have direct access to AWS. Instead, the application communicates with an identity provider (IdP) to authenticate the user. The IdP gets the user information from your organization’s identity store (such as an LDAP directory) and then generates a SAML assertion that includes authentication and authorization information about that user. The application then uses that assertion to make a call to the AssumeRoleWithSAML API to get temporary security credentials. The app can then use those credentials to access a folder in the S3 bucket that’s specific to the user. http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_saml.html

NEW QUESTION 38
You are designing a multi-platform web application for AWS. The application will run on EC2 instances and will be accessed from PCs. tablets and smart phones

Supported accessing platforms are Windows. MACOS. IOS and Android Separate sticky session and SSL certificate setups are required for different platform types which of the following describes the most cost effective and performance efficient architecture setup?

A. Setup a hybrid architecture to handle session state and SSL certificates on-prem and separate EC2 Instance groups running web applications for different platform types running in a VPC
B. Set up one ELB for all platforms to distribute load among multiple instance under it Each EC2 instance implements ail functionality for a particular platform.
C. Set up two ELBs The first ELB handles SSL certificates for all platforms and the second ELB handles session stickiness for all platforms for each ELB run separate EC2 instance groups to handle the web application for each platform.
D. Assign multiple ELBS to an EC2 instance or group of EC2 instances running the common components of the web application, one ELB for each platform type Session stickiness and SSL termination are done at the ELBs.

Answer: D

Explanation:
One ELB cannot handle different SSL certificates but since we are using sticky sessions it must be handled at the ELB level. SSL could be handled on the EC2 instances only with TCP configured ELB,
ELB supports sticky sessions only in HTTP/HTTPS configurations.
The way the Elastic Load Balancer does session stickiness is on a HTTP/HTTPS listener is by utilizing an HTTP cookie. If SSL traffic is not terminated on the Elastic Load Balancer and is
terminated on the back-end instance, the Elastic Load Balancer has no visibility into the HTTP headers and therefore cannot set or read any of the HTTP headers being passed back and forth. http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-sticky-sessions.html

NEW QUESTION 39
Your company produces customer commissioned one-of-a-kind skiing helmets combining nigh fashion with custom technical enhancements Customers can show off their Individuality on the ski slopes and have access to head-up-displays. GPS rear-view cams and any other technical innovation they wish to embed in the helmet.
The current manufacturing process is data rich and complex including assessments to ensure that the custom electronics and materials used to assemble the helmets are to the highest standards Assessments are a mixture of human and automated assessments you need to add a new set of
assessment to model the failure modes of the custom electronics using GPUs with CUDA, across a cluster of servers with low latency networking.
What architecture would allow you to automate the existing process using a hybrid approach and ensure that the architecture can support the evolution of processes over time?

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

Answer: B

NEW QUESTION 40
Your department creates regular analytics reports from your company's log files. All log data is collected in Amazon S3 and processed by daily Amazon Elastic MapReduce (EMR) jobs that generate daily PDF reports and aggregated tables in CSV format for an Amazon Redshift data warehouse. Your CFO requests that you optimize the cost structure for this system.
Which of the following alternatives will lower costs without compromising average performance of the system or data integrity for the raw data?

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

Answer: D

Explanation:
Reserved Instances (a.k.a. Reserved Nodes) are appropriate for steady-state production workloads, and offer significant discounts over On-Demand pricing. https://aws.amazon.com/redshift
Q: What are some EMR best practices?
If you are running EMR in production you should specify an AMI version, Hive version, Pig version, etc. to make sure the version does not change unexpectedly (e.g. when EMR later adds support for a newer version). If your cluster is mission critical, only use Spot instances for task nodes because if the Spot price increases you may lose the instances. In development, use logging and enable debugging
to spot and correct errors faster. If you are using GZIP, keep your file size to 1–2 GB because GZIP files cannot be split. Click here to download the white paper on Amazon EMR best practices. https://aws.amazon.com/elasticmapreduce/faqs

NEW QUESTION 45
Your website is serving on-demand training videos to your workforce. Videos are uploaded monthly in high resolution MP4 format. Your workforce is distributed globally often on the move and using
company-provided tablets that require the HTTP Live Streaming (HLS) protocol to watch a video. Your company has no video transcoding expertise and it required you may need to pay for a consultant.
How do you implement the most cost-efficient architecture without compromising high availability and quality of video delivery'?

A. A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queu
B. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few day
C. CloudFront to serve HLS transcoded videos from EC2.
D. Elastic Transcoder to transcode original high-resolution MP4 videos to HL
E. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few day

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

Answer: C

NEW QUESTION 50
You've been hired to enhance the overall security posture for a very large e-commerce site. They
have a well architected multi-tier application running in a VPC that uses ELBs in front of both the web and the app tier with static assets served directly from S3 They are using a combination of RDS and DynamoOB for their dynamic data and then archiving nightly into S3 for further processing with EMR They are concerned because they found questionable log entries and suspect someone is attempting
to gain unauthorized access.
Which approach provides a cost effective scalable mitigation to this kind of attack?

A. Recommend that they lease space at a DirectConnect partner location and establish a 1GDirectConnect connection to theirvPC they would then establish Internet connectivity into theirspace, filter the traffic in hardware Web Application Firewall (WAF). And then pass the traffic through the DirectConnect connection into their application running in their VPC,
B. Add previously identified hostile source IPs as an explicit INBOUND DENY NACL to the web tier subnet.
C. Add a WAF tier by creating a new ELB and an AutoScalmg group of EC2 Instances running a hostbased WAF They would redirect Route 53 to resolve to the new WAF tier ELB The WAF tier wouldthier pass the traffic to the current web tier The web tier Security Groups would be updated to only allow traffic from the WAF tier Security Group
D. Remove all but TLS 1 2 from the web tier ELB and enable Advanced Protocol Filtering This will enable the ELB itself to perform WAF functionality.

Answer: C

NEW QUESTION 52
An administrator is using Amazon CloudFormation to deploy a three tier web application that consists of a web tier and application tier that will utilize Amazon DynamoDB for storage when creating the CloudFormation template which of the following would allow the application instance access to the DynamoDB tables without exposing API credentials?

A. Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and associate the Role to the application instances by referencing an instance profile.
B. Use the Parameter section in the Cloud Formation template to 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 CloudFormation template that has permissions to read and write from the required DynamoDB table, use the GetAtt function to retrieve the Access and secret keys and pass them to the application instance through user-data.

Answer: C

NEW QUESTION 57
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 S3 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 S3 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 instanceid 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 certificateand 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 a certificate signature (hat contains the specific instance-id.

Answer: A

Explanation:
http://jayendrapatil.com/tag/iam/

NEW QUESTION 59
You are designing an SSUTLS solution that requires HTTPS clients to be authenticated by the Web server using client certificate authentication. The solution must be resilient.
Which of the following options would you consider for configuring the web server infrastructure? (Choose two.)

A. Configure ELB with TCP listeners on TCP/4d3. And place the Web servers behind it.
B. Configure your Web servers with EIPS Place the Web servers in a Route53 Record Set and configure health checks against all Web servers.
C. Configure ELB with HTTPS listeners, and place the Web servers behind it.
D. Configure your web servers as the origins for a CloudFront distributio
E. Use custom SSL certificateson your CloudFront distributio

Answer: AB

Explanation:
This question is regarding “two-way” SSL authentication.
Currently, ELBs cannot support authentication for the client side SSL/TLS cert required for two-way SSL authentication to succeed. Therefore, you only have two options:
\A. Configure the ELB with a TCP/443 listener. This is effectively TLS “pass through” mode, where the TLS connection does not terminate on the ELB, it is passed through and decrypted on the back-end servers. This will cause quite a bit of CPU overhead on the back-end instances, due to the lack of TLS offload that cannot happen on the ELB, so an auto-scaling group which monitors the web server CPU metrics would be essential here. (Not that you probably wouldn’t have it anyway, just saying!)
\B. Don’t use an ELB. Just have the web servers act as the endpoint for the traffic, and let Route53 DNS serve in the place of the ELB by load balancing client DNS queries across the web servers. C and D are not options here, since neither are supported by AWS.

NEW QUESTION 62
You have an application running on an EC2 Instance which will allow users to download files from a private S3 bucket using a pre-assigned URL. Before generating the URL the application should verify the existence of the file in S3.
How should the application use AWS credentials to access the S3 bucket securely?

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

Answer: C

Explanation:
Reference
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-metadata.html

NEW QUESTION 65
You are designing a social media site and are considering how to mitigate distributed denial-ofservice (DDoS) attacks. Which of the below are viable mitigation techniques? (Choose three.)

A. Add multiple elastic network interfaces (ENIs) to each EC2 instance to increase the network bandwidth.
B. Use dedicated instances to ensure that each instance has the maximum performance possible.
C. Use an Amazon CloudFront distribution for both static and dynamic content.
D. Use an Elastic Load Balancer with auto scaling groups at the we
E. App and Amazon Relational Database Service (RDS) tiers
F. Add alert Amazon CloudWatch to look for high Network in and CPU utilization.
G. Create processes and capabilities to quickly add and remove rules to the instance OS firewal

Answer: CDE

NEW QUESTION 68
Your fortune 500 company has under taken a TCO analysis evaluating the use of Amazon S3 versus acquiring more hardware. The outcome was that all employees would be granted access to use Amazon S3 for storage of their personal documents.
Which of the following will you need to consider so you can set up a solution that incorporates single sign-on from your corporate AD or LDAP directory and restricts access for each user to a designated user folder in a bucket? (Choose three.)

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

Answer: ABD

NEW QUESTION 71
You have a periodic Image analysis application that gets some files. In Input analyzes them and for each file writes some data in output to a 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. S3 to store I/O file
B. SQS 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 SQS queue
D. EBS with Provisioned IOPS (PIOPS) to store I/O file
E. SNS to distribute elaboration commands to a group of hosts working in parallel Auto Scaling to dynamically size the group of hosts depending on the number of SNS notifications
F. S3 to store I/O files, SNS to distribute 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 I/O 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 SQS queue.

Answer: D

Explanation:
Amazon EBS allows you to create storage volumes and attach them to Amazon EC2 instances. Once attached, you can create a file system on top of these

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

NEW QUESTION 73
Your company is getting ready to do a major public announcement of a social media site on AWS. The website is running on EC2 instances deployed across multiple Availability Zones with a Multi-AZ RDS MySQL Extra Large DB Instance. The site performs a high number of small reads and writes per second and relies on an eventual consistency model. After comprehensive tests you discover that
there is read contention on RDS MySQL. Which are the best approaches to meet these requirements? (Choose two.)

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

Answer: AD

NEW QUESTION 76
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 two.)

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 the 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 81
To serve Web traffic for a popular product your chief financial officer and IT director have purchased 10 ml large heavy utilization Reserved Instances (RIs) evenly spread across two availability zones: Route 53 is used to deliver the traffic to an Elastic Load Balancer (ELB). After several months, the product grows even more popular and you need additional capacity. As a result, your company purchases two C3.2xlarge medium utilization Ris.
You register the two c3 2xlarge instances with your ELB and quickly find that the ml large instances are at 100% of capacity and the c3 2xlarge instances have significant capacity that's unused.
Which option is the most cost effective and uses EC2 capacity most effectively?

A. Use a separate ELB for each instance type and distribute load to ELBs with Route 53 weighted round robin
B. Configure Autoscaning group and Launch Configuration with ELB to add up to 10 more on-demand mi large instances when triggered by Cloudwatch shut off c3 2xiarge instances
C. Route traffic to EC2 ml large and c3 2xlarge instances directly using Route 53 latency based routing and health checks shut off ELB
D. Configure ELB with two c3 2xlarge Instances and use on-demand Autoscailng group for up to two additional c3.2xlarge instances Shut on mi .large instances.

Answer: A

Explanation:
Weighted Routing Policy
Use the weighted routing policy when you have multiple resources that perform the same function (for example, web servers that serve the same website) and you want Amazon Route 53 to route traffic to those resources in proportions that you specify (for example, one quarter to one server and three quarters to the other).
For more information about weighted resource record sets, see Weighted Routing.

NEW QUESTION 86
A large real-estate brokerage is exploring the option or adding a cost-effective location based alert to their existing mobile application. The application backend infrastructure currently runs on AWS Users who opt in to this service will receive alerts on their mobile device regarding real-estate otters in proximity to their location. For the alerts to be relevant delivery time needs to be in the low minute count the existing mobile app has 5 million users across the us. Which one of the following architectural suggestions would you make to the customer?

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

Answer: A

Explanation:
AWS using SQS to store the message from mobile apps,and using AWS Mobile Push to send offers to mobile apps.

NEW QUESTION 90
You are developing a new mobile application and are considering storing user preferences in AWS.2w. This would provide a more uniform cross-device experience to users using multiple mobile devices to access the application. The preference data for each user is estimated to be 50KB in size Additionally 5 million customers are expected to use the application on a regular basis. The solution needs to be cost-effective, highly available, scalable and secure, how would you design a solution to meet the above requirements?

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

Answer: B

Explanation:
https://aws.amazon.com/blogs/aws/fine-grained-access-control-for-amazon-dynamodb/ Here are some of the things that you can build using fine-grained access control:
A mobile app that displays information for nearby airports, based on the user's location. The app can access and display attributes such airline names, arrival times, and flight numbers. However,
it cannot access or display pilot names or passenger counts.
A mobile game which stores high scores for all users in a single table. Each user can update their own scores, but has no access to the other ones.

NEW QUESTION 94
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 accoun
E. Create cross-account roles in the Dev and Test accounts that have full Admin permissions and grant the Master account access.
F. Link the accounts using Consolidated Billin
G. This will give IAM users in the Master account access to resources in the Dev and Test accounts

Answer: C

NEW QUESTION 99
A web company is looking to implement an intrusion detection and prevention system into their deployed VPC. This platform should have the ability to scale to thousands of instances running inside of the VPC, How should they architect their solution to achieve these goals?

A. Configure an instance with monitoring software and the elastic network interface (ENI) set to promiscuous mode packet sniffing to see an traffic across the VPC,
B. Create a second VPC and route all traffic from the primary application VPC through the second VPC where the scalable virtualized IDS/IPS platform resides.
C. Configure servers running in the VPC using the host-based 'route' commands to send all traffic through the platform to a scalable virtualized IDS/IPS.
D. Configure each host with an agent that collects all network traffic and sends that traffic to the IDS/IPS platform for inspection.

Answer: D

Explanation:
\A. Not possible to set an instance's NIC into promiscuous mode.
\B. Incorrect… VPC peering connections are not "transitive", i.e. you cannot pass traffic through a VPC peering connection into another VPC, and then have that other VPC send the traffic to some third VPC, or the Internet, or a VPN, or a direct connect circuit. (I would assume AWS does not allow redistribution of routes from one VPC's back-end VRF into another VPC's back-end VRF, unless it is that first VPC's CIDR block? Someone from AWS would have to chime in here, and they're probably not going to tell us.)
\C. This one is incorrect because adding static routes on an instance won't affect the routing from any point after the packet leaves the instance's NIC. AWS will check the destination IP address in the packet header and forward according to the VPC routing table's routes. You'd need to make routing changes in the VPC route table for that instance's traffic to get sent through another device (e.g. NAT gateway, VPN instance, or security proxy in this case). (You could tunnel/proxy the traffic over through the IPS tier by changing the destination IP address in the IP header of the packet before it left the instance. But choice C did not state anything about doing anything like that. It just said add a static route on the instance, which does not change the destination IP address in the IP header of the packet.)
\D. Correct, this is the standard approach, and is definitely scalable.

NEW QUESTION 104
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 AMI 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: B

Explanation:
You only need one stack to contain two layers:
– one layer for the Java/Tomcat instances
– one layer for DynamoDB
You'd only need one custom recipe because the only OpsWorks Lifecycle Event that would be involved in rolling out the new chat feature would be "Deploy". (Or you could implement it in "Setup" if you choose to make including the chat app a new baseline standard for your instances in that layer. But even then, you'd only have one custom recipe because there would be no need to customize the "Deploy" event to install the chat app if you already installed out the chat app in "Setup".) So you'd need a custom recipe for that one lifecycle event. And it would only be used for the "Deploy" lifecycle event on the app layer, not on the DB layer

NEW QUESTION 106
Your firm has uploaded a large amount of aerial image data to S3. In the past, in your on-premises environment, you used a dedicated group of servers to oaten process this data and used Rabbit MQ - An open source messaging system to get job information to the servers. Once processed the data would go to tape and be shipped offsite. Your manager told you to stay with the current design, and leverage AWS archival storage and messaging services to minimize cost. Which is correct?

A. Use SQS for passing job messages use Cloud Watch alarms to terminate EC2 worker instances when they become idl
B. Once data is processed, change the storage class of the S3 objects to Reduced Redundancy Storage.
C. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SOS Once data is processed,
D. Change the storage class of the S3 objects to Reduced Redundancy Storag
E. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SQS Once data is processed, change the storage class of the S3 objects to Glacier.
F. Use SNS to pass job messages use Cloud Watch alarms to terminate spot worker instances when they become idl
G. Once data is processed, change the storage class of the S3 object to Glacier.

Answer: C

Explanation:
The question key part to focus on is "and leverage AWS archival storage and messaging services to minimize cost."
For that the storage that is the lowest cost in the answers is Glacier, in addition, the messaging cost is less for SQS then for SNS if they both exceed 1 million transactions which is free. The
only answer that satisfies the above two criteria is answer C. Also, there does not seem to be an urgency in speed of messaging therefore SQS satisfies that need. SNS being more real time delivery mechanism.

NEW QUESTION 110
You must assign each server to at least _____ security group

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

Answer: D

Explanation:
Your AWS account automatically has a default security group per region for EC2-Classic. When you create a VPC, we automatically create a default security group for the VPC. If you don't specify a different security group when you launch an instance, the instance is automatically associated with the appropriate default security group. http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html

NEW QUESTION 111
Before I delete an EBS volume, what can I do if I want to recreate the volume later?

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

Answer: B

Explanation:
After you no longer need an Amazon EBS volume, you can delete it. After deletion, its data is gone and the volume can't be attached to any instance. However, before deletion, you can store a snapshot of the volume, which you can use to re-create the volume later.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-deleting-volume.html

NEW QUESTION 115
What does RRS stand for when talking about S3?

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

D. Reduced Redundancy Storage

Answer: D

Explanation:
In Amazon S3, RRS stands for Reduced Redundancy Storage. Reduced redundancy storage stores objects on multiple devices across multiple facilities, providing 400 times the durability of a typical disk drive, but it does not replicate objects as many times as Amazon S3 standard storage. In addition, reduced redundancy storage is designed to sustain the loss of data in a single facility. http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingRRS.html

NEW QUESTION 116
What does Amazon SWF stand for?

A. Simple Web Flow
B. Simple Work Flow
C. Simple Wireless Forms
D. Simple Web Form

Answer: B

NEW QUESTION 118
What does Amazon EC2 provide?

A. Virtual servers in the Cloud.
B. A platform to run code (Java, PHP, Python), paying on an hourly basis.
C. Computer Clusters in the Cloud.
D. Physical servers, remotely managed by the custome

Answer: A

NEW QUESTION 122
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

Explanation:
When you view the block device mapping for your instance, you can see only the EBS volumes, not the instance store volumes. You can use instance metadata to query the complete block device mapping. The base URI for all requests for instance metadata is http://169.254.169.254/latest/. http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/block-device-mappingconcepts. html#bdm-instance-metadata

NEW QUESTION 123
What does Amazon Elastic Beanstalk provide?

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

Answer: B

NEW QUESTION 128
What will be the status of the snapshot until the snapshot is complete.

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

Answer: D

Explanation:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html

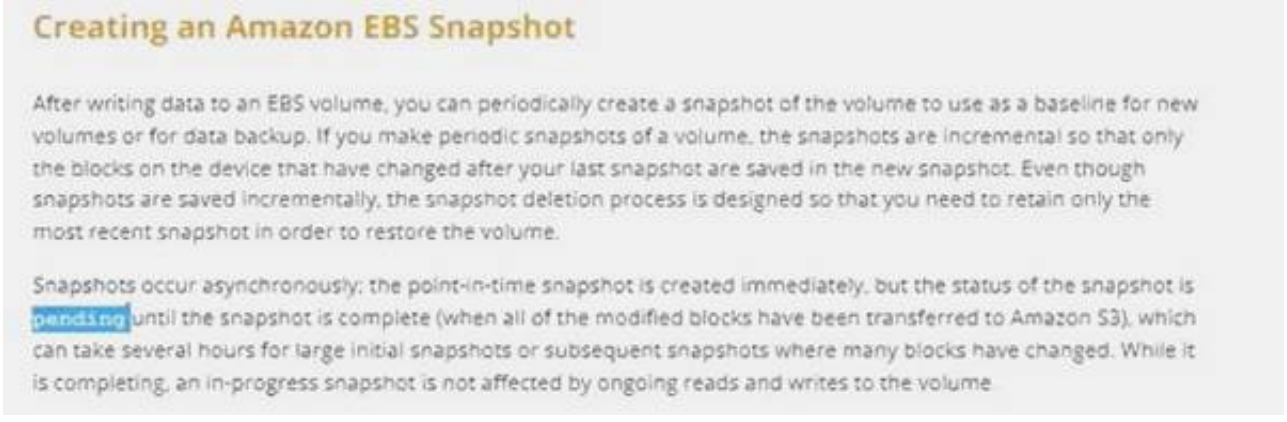
Creating an Amazon EBS Snapshot

After writing data to an EBS volume, you can periodically create a snapshot of the volume to use as a baseline for new volumes or for data backup. If you make periodic snapshots of a volume, the snapshots are incremental so that only the blocks on the device that have changed after your last snapshot are saved in the new snapshot. Even though snapshots are saved incrementally, the snapshot deletion process is designed so that you need to retain only the most recent snapshot in order to restore the volume.

Snapshots occur asynchronously; the point-in-time snapshot is created immediately, but the status of the snapshot is pending until the snapshot is complete (when all of the modified blocks have been transferred to Amazon S3), which can take several hours for large initial snapshots or subsequent snapshots where many blocks have changed. While it is completing, an in-progress snapshot is not affected by ongoing reads and writes to the volume

NEW QUESTION 130
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 132
Amazon RDS automated backups and DB Snapshots are currently supported for only the _ storage engine

A. InnoDB
B. MyISAM

Answer: A

NEW QUESTION 134
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 138
What is the maximum key length of a tag?

A. 512 Unicode characters
B. 64 Unicode characters
C. 256 Unicode characters
D. 128 Unicode characters

Answer: D

NEW QUESTION 143
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

Explanation:
Changes to the backup window take effect immediately, with the limitations that the specified backup window must be at least 10 minutes in the future, and the backup window cannot overlap with the weekly maintenance window for the instance.

NEW QUESTION 144
Using Amazon CloudWatch's Free Tier, what is the frequency of metric updates which you receive?

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

Answer: A

Explanation:
You can get started with Amazon CloudWatch for free. Many applications should be able to operate within these free tier limits.
New and existing customers also receive 3 dashboards of up to 50 metrics each per month at no
additional charge Basic Monitoring metrics (at five-minute frequency) for Amazon EC2 instances are free of charge, as are all metrics for Amazon EBS volumes, Elastic Load Balancers, and Amazon RDS DB instances.
https://aws.amazon.com/cloudwatch/pricing/

NEW QUESTION 147
Which is the default region in AWS?

A. eu-west-1
B. us-east-1
C. us-east-2

D. ap-southeast-1

Answer: B

Explanation:
All the main AWS services (except Route 53 & CloudFront) allow you to select which region you would like to use. The US East (N. Virginia) is the default region. You can change the region by using the dropdown menu in the top right of the management console.

NEW QUESTION 149
What are the two types of licensing options available for using Amazon RDS for Oracle?

A. BYOL and Enterprise License
B. BYOL and License Included
C. Enterprise License and License Included
D. Role based License and License Included

Answer: B

Explanation:
https://aws.amazon.com/rds/oracle/

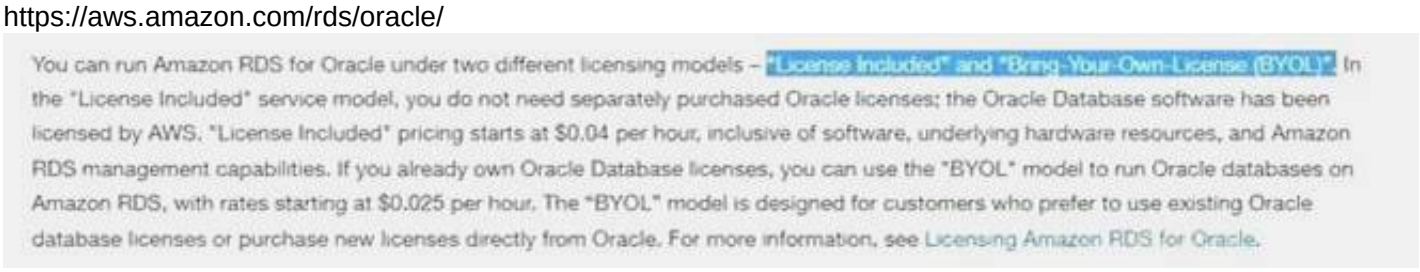
You can run Amazon RDS for Oracle under two different licensing models – "License Included" and "Bring-Your-Own-License (BYOL)". In the "License Included" service model, you do not need separately purchased Oracle licenses; the Oracle Database software has been licensed by AWS. "License Included" pricing starts at $0.04 per hour, inclusive of software, underlying hardware resources, and Amazon RDS management capabilities. If you already own Oracle Database licenses, you can use the "BYOL" model to run Oracle databases on Amazon RDS, with rates starting at $0.025 per hour. The "BYOL" model is designed for customers who prefer to use existing Oracle database licenses or purchase new licenses directly from Oracle. For more information, see Licensing Amazon RDS for Oracle.

NEW QUESTION 152
What does a "Domain" refer to in Amazon SWF?

A. A security group in which only tasks inside can communicate with each other
B. A special type of worker
C. A collection of related Workflows
D. The DNS record for the Amazon SWF service

Answer: C

Explanation:
Domains provide a way of scoping Amazon SWF resources within your AWS account. All the components of a workflow, such as the workflow type and activity types, must be specified to be in a domain. It is possible to have more than one workflow in a domain; however, workflows in different domains cannot interact with each other. http://docs.aws.amazon.com/amazonswf/latest/developerguide/swf-dev-domain.html

NEW QUESTION 153
Disabling automated backups disable the point-in-time recovery.

A. if configured to can
B. will never
C. will

Answer: C

NEW QUESTION 158
Out of the stripping options available for the EBS volumes, which one has the following disadvantage: 'Doubles the amount of I/O required from the instance to EBS compared to RAID 0, because you're mirroring all writes to a pair of volumes, limiting how much you can stripe.' ?

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

Answer: B

Explanation:
http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/raid-config.html raid 0 and 1 are the common types. Raid 5 and 6 are not recommended because of the extended stripe. If you encounter this question on the exam I suspect the answer options will be different.
Raid 1 Disadvantage
Does not provide a write performance improvement; requires more Amazon EC2 to Amazon EBS bandwidth than non- RAID configurations because the data is written to multiple volumes simultaneously.
Raid 0 Disadvantage
Performance of the stripe is limited to the worst performing volume in the set. Loss of a single volume results in a complete data loss for the array.
Raid 5 and 6 notes
RAID 5 and RAID 6 are not recommended for Amazon EBS because the parity write operations of these RAID modes consume some of the IOPS available to your volumes. Depending on the
configuration of your RAID array, these RAID modes provide 20-30% fewer usable IOPS than a RAID 0 configuration. Increased cost is a factor with these RAID modes as well; when using identical volume sizes and speeds, a 2-volume RAID 0 array can outperform a 4-volume RAID 6 array that costs twice
as much.

NEW QUESTION 159
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

Explanation:
https://aws.amazon.com/rds/faqs/

Q: Can I create a Read Replica of another Read Replica?

Amazon RDS for MySQL and *MariaDB*: You can create a second-tier Read Replica from an existing first-tier Read Replica. By creating a second-tier Read Replica, you may be able to move some of the replication load from the master database instance to a first-tier Read Replica. Please note that a second-tier Read Replica may lag further behind the master because of additional replication latency introduced as transactions are replicated from the master to the first tier replica and then to the second-tier replica.

Amazon RDS for PostgreSQL: Read Replicas of Read Replicas are not currently supported.

NEW QUESTION 162
Can Amazon S3 uploads resume on failure or do they need to restart?

A. Restart from beginning
B. You can resume them, if you flag the "resume on failure" option before uploading.
C. Resume on failure
D. Depends on the file size

Answer: C

NEW QUESTION 165
True or False: Manually created DB Snapshots are deleted after the DB Instance is deleted.

A. TRUE
B. FALSE

Answer: B

Explanation:
If you choose not to create a final DB snapshot, you will not be able to later restore the DB instance to its final state. When you delete a DB instance, all automated backups are deleted and cannot be recovered. Manual DB snapshots of the instance are not deleted.
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_DeleteInstance.html

NEW QUESTION 170
What is Amazon Glacier?

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

Answer: C

Explanation:
Amazon Glacier is an extremely low-cost storage service that provides durable storage with security features for data archiving and backup.

NEW QUESTION 171
What is the durability of S3 RRS?

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

Answer: A

Explanation:
RRS = Reduced Redundancy Storage

	Standard	Standard - Infrequent Access	Reduced Redundancy Storage
Durability	99.999999999%	99.999999999%	99.99%

NEW QUESTION 172
Is Federated Storage Engine currently supported by Amazon RDS for MySQL?

A. Only for Oracle RDS instances
B. No
C. Yes
D. Only in VPC

Answer: B

NEW QUESTION 175
Is there a limit to how many groups a user can be in?

A. Yes for all users
B. Yes for all users except root
C. No
D. Yes unless special permission granted

Answer: A

Explanation:
Currently you can request to increase the limit on users per AWS account, groups per AWS account, roles per AWS account, instance profiles per AWS account, and server certificates per
AWS account.
This never states “groups a user can be in”

NEW QUESTION 176
Will my standby RDS instance be in the same Availability Zone as my primary?

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

Answer: D

NEW QUESTION 179
What does the 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 group.

Answer: B

NEW QUESTION 182
What happens to the data on an instance if the instance reboots (intentionally or unintentionally)?

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

Answer: B

Explanation:
Instance Store Lifetime
You can specify instance store volumes for an instance only when you launch it. The data in an instance store persists only during the lifetime of its associated instance. If an instance reboots (intentionally or unintentionally), data in the instance store persists. However, data in the instance store is lost under the following circumstances:
The underlying disk drive fails The instance stops
The instance terminates http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html

NEW QUESTION 183
How many types of block devices does Amazon EC2 support?

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

Answer: A

Explanation:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/block-device-mapping-concepts.html Amazon EC2 supports two types of block devices:
Instance store volumes (virtual devices whose underlying hardware is physically attached to the host computer for the instance)
EBS volumes (remote storage devices)

A block device mapping defines the block devices (instance store volumes and EBS volumes) to attach to an instance.

Block Device Mapping Concepts

A *block device* is a storage device that moves data in sequences of bytes or bits (blocks). These devices support random access and generally use buffered I/O. Examples include hard disks, CD-ROM drives, and flash drives. A block device can be physically attached to a computer or accessed remotely as if it were physically attached to the computer. Amazon EC2 supports two types of block devices:

- Instance store volumes (virtual devices whose underlying hardware is physically attached to the host computer for the instance)
- EBS volumes (remote storage devices)

NEW QUESTION 185
IAM provides several policy templates you can use to automatically assign permissions to the groups you create. The ___ policy template gives the Admins group permission to access all account resources, except your AWS account information

A. Read Only Access
B. Power User Access
C. AWS Cloud Formation Read Only Access
D. Administrator Access

Answer: D

Explanation:
AWS managed policies are designed to provide permissions for many common use cases. For example, there are AWS managed policies that define typical permissions for administrators (all access), for power users (all access except IAM), and for other various levels of access to AWS services. AWS managed policies make it easier for you to assign appropriate permissions to users, groups, and roles than if you had to write the policies yourself.
http://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_managed-vs-inline.html

NEW QUESTION 189
IAM's Policy Evaluation Logic always starts with a default ____ for every request, except for those that use the AWS account's root security credentials b

A. Permit
B. Deny
C. Cancel

Answer: B

Explanation:
http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_evaluation-logic.html

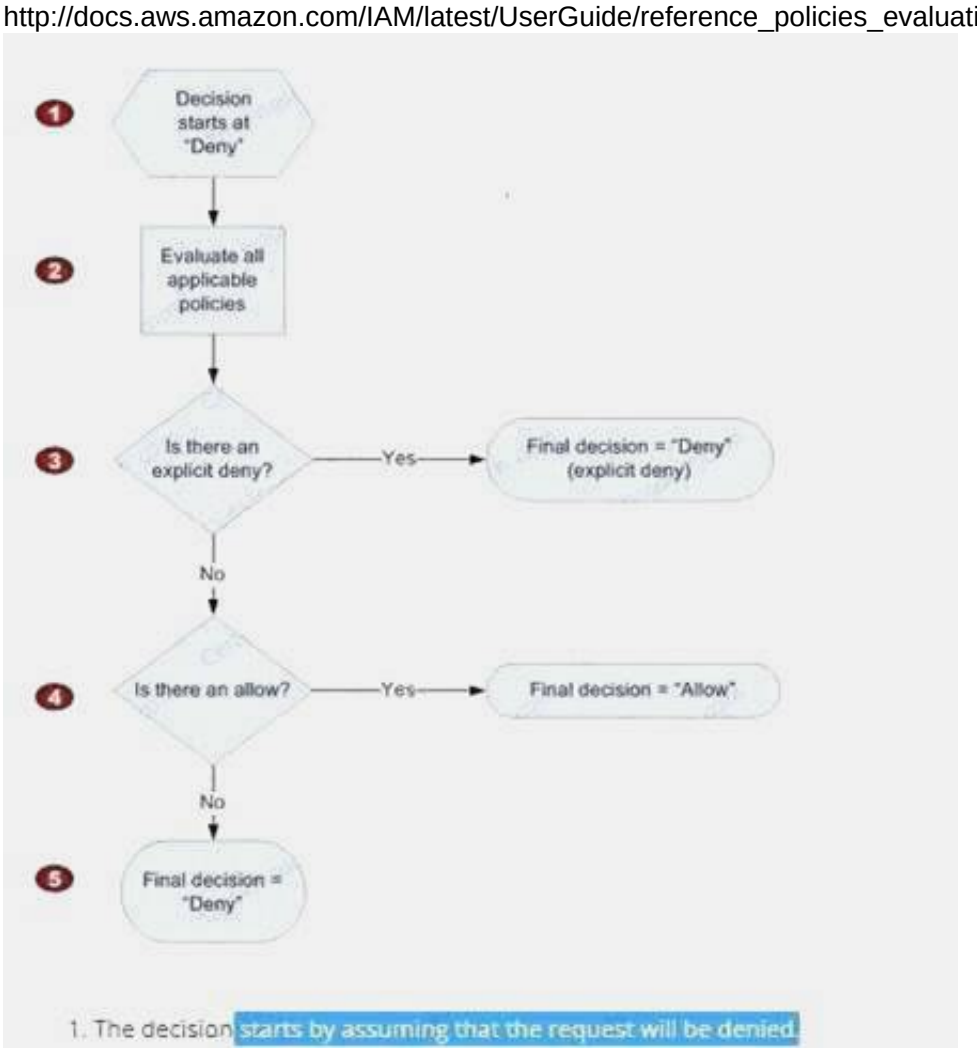

NEW QUESTION 193
What is Oracle SQL Developer?

A. An AWS developer who is an expert in Amazon RDS using both the Oracle and SQL Server DB engines
B. A graphical Java tool distributed without cost by Oracle.
C. It is a variant of the SQL Server Management Studio designed by Microsoft to support Oracle DBMS functionalities
D. A different DBMS released by Microsoft free of cost

Answer: B

Explanation:
http://www.oracle.com/technetwork/developer-tools/sql-developer/what-is-sqldev-093866.html

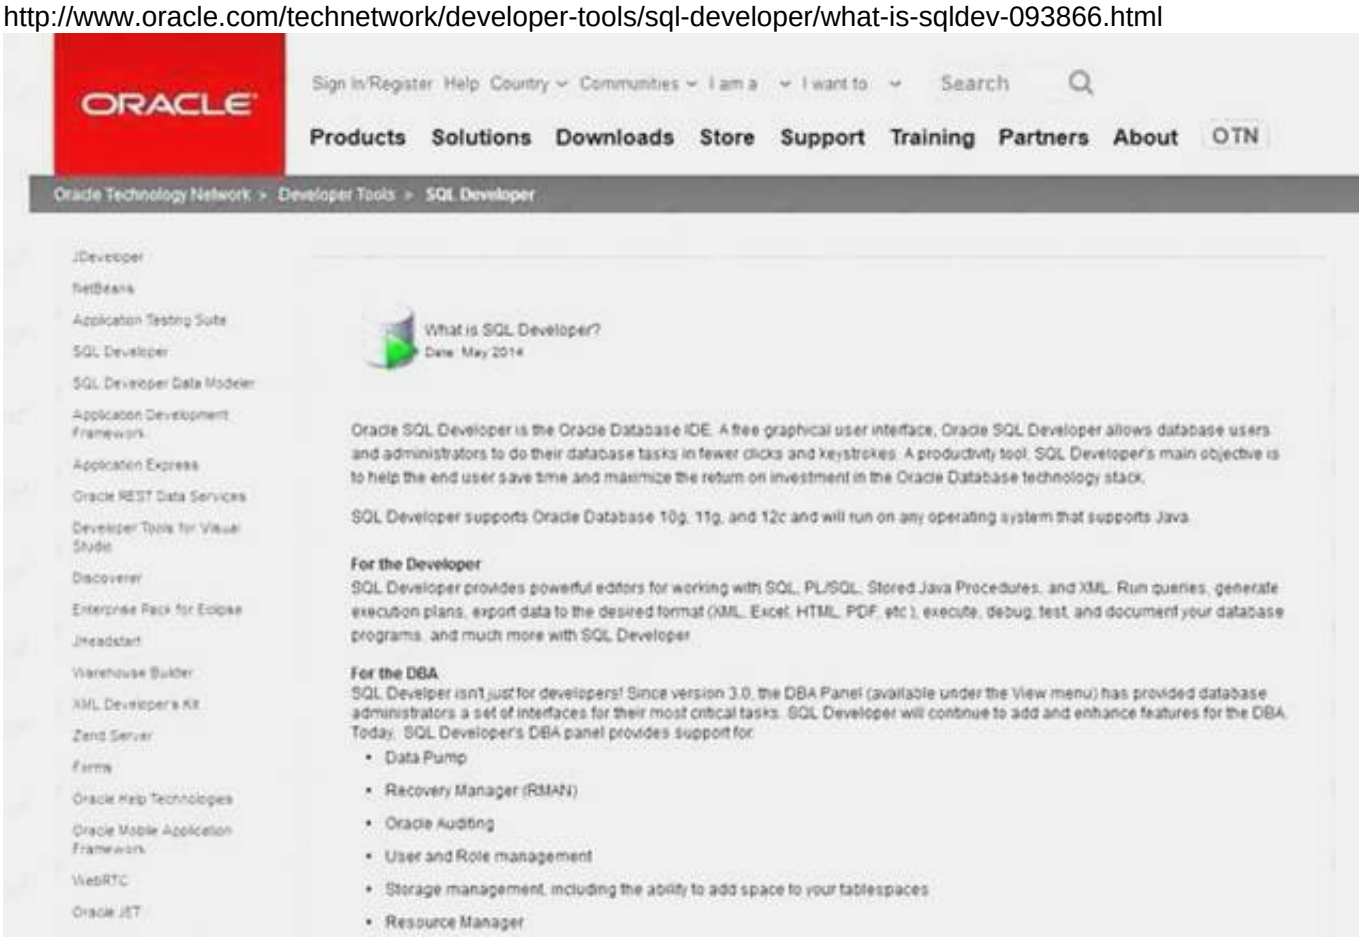

NEW QUESTION 194
To view information about an Amazon EBS volume, open the Amazon EC2 console at https://console.aws.amazon.com/ec2/, click _____ in the Navigation pane.

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

Answer: D

Explanation:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-describing-volumes.html

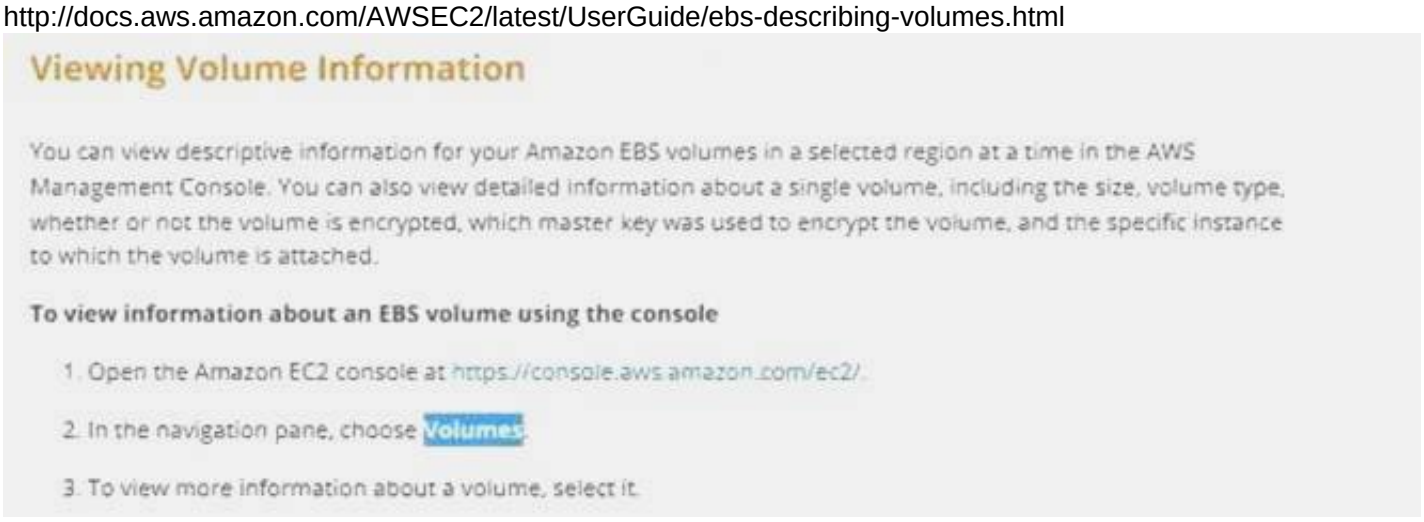

NEW QUESTION 196
Using Amazon IAM, can I give permission based on organizational groups?

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

Answer: C

Explanation:
An IAM group is a collection of IAM users. You can use groups to specify permissions for a collection of users, which can make those permissions easier to manage for those users. http://docs.aws.amazon.com/IAM/latest/UserGuide/id.html

NEW QUESTION 198
While creating the snapshots using the API, which Action should I be using?

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

Answer: D

Explanation:
http://docs.aws.amazon.com/AWSEC2/latest/CommandLineReference/ApiReference-cmd- CreateSnapshot.html

NEW QUESTION 201
What happens to the I/O operations while you take a database snapshot?

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

Answer: A

Explanation:
Creating this DB snapshot on a Single-AZ DB instance results in a brief I/O suspension that typically lasting no more than a few minutes. Multi-AZ DB instances are not affected by this I/O
suspension since the backup is taken on the standby.

NEW QUESTION 204
Read Replicas require a transactional storage engine and are only supported for the _____ storage engine

A. OracleISAM
B. MSSQLDB
C. InnoDB
D. MyISAM

Answer: C

Explanation:
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

- Using a non-transactional storage engine such as MyISAM. Read replicas require a transactional storage engine. Replication is only supported for the InnoDB storage engine on MySQL and the XtraDB storage engine on MariaDB.

NEW QUESTION 207
In the 'Detailed' monitoring data available for your Amazon EBS volumes, Provisioned IOPS volumes automatically send _____ minute metrics to Amazon CloudWatch.

A. 3
B. 1
C. 5
D. 2

Answer: B

NEW QUESTION 212
Which service enables AWS customers to manage users and permissions in AWS?

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

Answer: B

NEW QUESTION 214
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 215
What is the command line instruction for running the remote desktop client in Windows?

A. desk.cpl
B. mstsc

Answer: B

NEW QUESTION 218
Amazon RDS automated backups and DB Snapshots are currently supported for only the _____ storage engine

A. MyISAM
B. InnoDB

Answer: B

NEW QUESTION 220
In regards to IAM you can edit user properties later, but you cannot use the console to change the ____.

A. user name
B. password
C. default group

Answer: A

NEW QUESTION 223
Can I use Provisioned IOPS with VPC?

A. Only Oracle based RDS
B. No
C. Only with MSSQL based RDS
D. Yes for all RDS instances

Answer: D

NEW QUESTION 228
Fill in the blanks: "To ensure failover capabilities, consider using a ____ for incoming traffic on a network interface".

A. primary public IP
B. secondary private IP
C. secondary public IP
D. add on secondary IP

Answer: B

Explanation:
To ensure failover capabilities, consider using a secondary private IP for incoming traffic on an elastic network interface. In the event of an instance failure, you can move the interface and/or secondary private IP address to a standby instance

NEW QUESTION 231
If I have multiple Read Replicas for my master DB Instance and I promote one of them, what happens to the rest of the Read Replicas?

A. The remaining Read Replicas will still replicate from the older master DB Instance
B. The remaining Read Replicas will be deleted
C. The remaining Read Replicas will be combined to one read replica

Answer: A

Explanation:
If a source DB instance has several Read Replicas, promoting one of the Read Replicas to a DB instance has no effect on the other replicas.

NEW QUESTION 236
What can I access by visiting the URL: http://status.aws.amazon.com/?

A. Amazon Cloud Watch
B. Status of the Amazon RDS DB
C. AWS Service Health Dashboard
D. AWS Cloud Monitor

Answer: C

NEW QUESTION 240
Please select the Amazon EC2 resource which cannot be tagged.

A. images (AMIs, kernels, RAM disks)
B. Amazon EBS volumes
C. Elastic IP addresses
D. VPCs

Answer: C

Explanation:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html#tag-restrictions

Resource	Tagging support	Tagging restrictions
AMI	Yes	None
Bundle task	No	
Customer gateway	Yes	None
Dedicated Host	No	
DHCP option	Yes	None
EBS volume	Yes	None
Instance store volume	No	
Elastic IP	No	
Egress-only Internet gateway	No	
Instance	Yes	None
Internet gateway	Yes	None
Key pair	No	
NAT gateway	No	
Network ACL	Yes	None
Network interface	Yes	None
Placement group	No	
Reserved Instance	Yes	None

NEW QUESTION 241
Because of the extensibility limitations of striped storage attached to Windows Server, Amazon RDS does not currently support increasing storage on a ____ DB Instance.

A. SQL Server
B. MySQL
C. Oracle

Answer: A

NEW QUESTION 245
How are the EBS snapshots saved on Amazon S3?

A. Exponentially
B. Incrementally
C. EBS snapshots are not stored in the Amazon S3
D. Decrementally

Answer: B

NEW QUESTION 246
Can I delete a snapshot of the root device of an EBS volume used by a registered AMI?

A. Only via API
B. Only via Console
C. Yes
D. No

Answer: D

Explanation:
Note that you can't delete a snapshot of the root device of an EBS volume used by a registered AMI. You must first deregister the AMI before you can delete the snapshot.
Source: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-deleting-snapshot.html

NEW QUESTION 249
The ____ service is targeted at organizations with multiple users or systems that use AWS products such as Amazon EC2, Amazon SimpleDB, and the AWS Management Console.

A. Amazon RDS
B. AWS Integrity Management
C. AWS Identity and Access Management
D. Amazon EMR

Answer: C

Explanation:
https://aws.amazon.com/documentation/iam/?nc1=h_ls

NEW QUESTION 254
True or False: Without IAM, you cannot control the tasks a particular user or system can do and what AWS resources they might use.

A. FALSE
B. TRUE

Answer: B

Explanation:
http://docs.aws.amazon.com/IAM/latest/UserGuide/getting-setup.html

NEW QUESTION 258
What is the default maximum number of MFA devices in use per AWS account (at the root account level)?

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

Answer: A

Explanation:
http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_iam-limits.html

NEW QUESTION 261
Do the Amazon EBS volumes persist independently from the running life of an Amazon EC2 instance?

A. Only if instructed to when created
B. Yes
C. No

Answer: B

Explanation:
Data persistence
An EBS volume is off-instance storage that can persist independently from the life of an instance. You continue to pay for the volume usage as long as the data persists.
References:

NEW QUESTION 262
What does Amazon ElastiCache provide?

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

Answer: C

NEW QUESTION 265
The one-time payment for Reserved Instances is ____ refundable if the reservation is cancelled.

A. always
B. in some circumstances
C. never

Answer: C

Explanation:
the one-time fee is non-refundable.
https://aws.amazon.com/ec2/purchasing-options/reserved-instances/buyer/

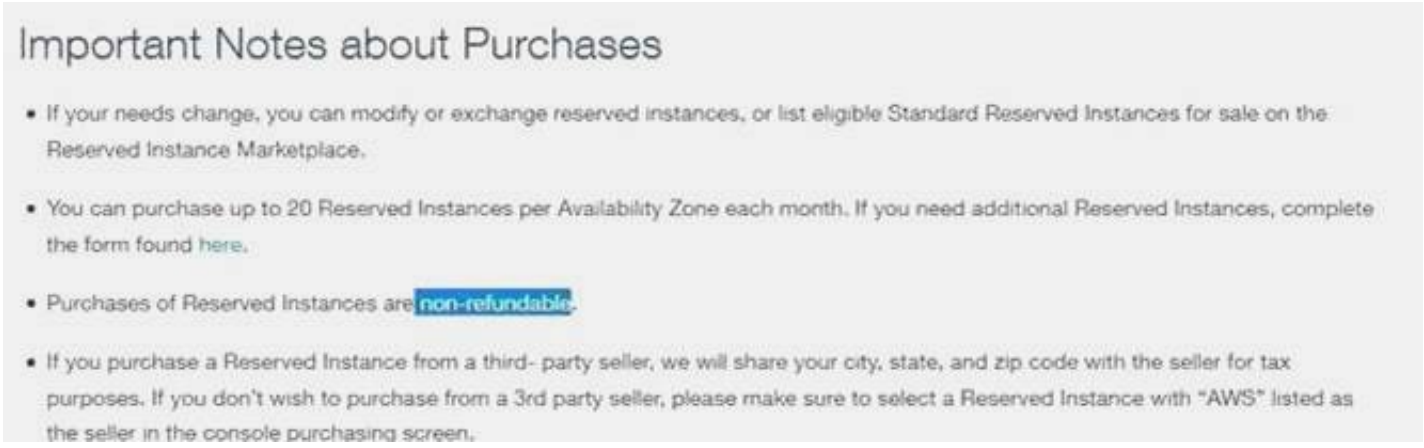

NEW QUESTION 267
Please select the Amazon EC2 resource which can be tagged.

A. key pairs
B. Elastic IP addresses
C. placement groups
D. Amazon EBS snapshots

Answer: C

Explanation:
Placement group and Elastic IP cannot be tagged. http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html Snapshots can be tagged: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html

Resource	Tagging support	Tagging restrictions
AMI	Yes	None
Bundle task	No	
Customer gateway	Yes	None
Dedicated Host	No	
DHCP option	Yes	None
EBS volume	Yes	None
Instance store volume	No	
Elastic IP	No	
Egress-only Internet gateway	No	
Instance	Yes	None
Internet gateway	Yes	None
Key pair	No	
NAT gateway	No	
Network ACL	Yes	None
Network interface	Yes	None
Placement group	No	
Reserved Instance	Yes	None
Reserved Instance listing	No	
Route table	Yes	None
Spot instance request	Yes	None
Security group - EC2-Classic	Yes	None
Security group - VPC	Yes	None
Snapshot	Yes	None

NEW QUESTION 269
Which of the following statements are true about Amazon Route 53 resource records? (Choose two.)

A. An Alias record can map one DNS name to another Amazon Route 53 DNS name.
B. A CNAME record can be created for your zone apex.
C. An Amazon Route 53 CNAME record can point to any DNS record hosted anywhere.
D. TTL can be set for an Alias record in Amazon Route 53.
E. An Amazon Route 53 Alias record can point to any DNS record hosted anywher

Answer: AC

NEW QUESTION 274
Amazon EC2 has no Amazon Resource Names (ARNs) because you can't specify a particular Amazon EC2 resource in an IAM policy.

A. TRUE
B. FALSE

Answer: B

Explanation:
http://blogs.aws.amazon.com/security/post/Tx29HCT3ABL7LP3/Resource-level-Permissions-for-EC2- Controlling- Management-Access-on-Specific-Ins

NEW QUESTION 279
Is the encryption of connections between my application and my DB Instance using SSL for the MySQL server engines available?

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

Answer: A

Explanation:
https://aws.amazon.com/rds/faqs/

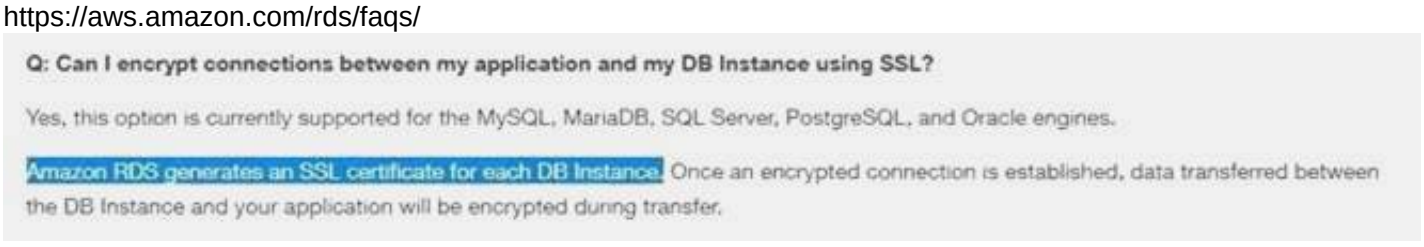
Q: Can I encrypt connections between my application and my DB Instance using SSL?

Yes, this option is currently supported for the MySQL, MariaDB, SQL Server, PostgreSQL, and Oracle engines.

Amazon RDS generates an SSL certificate for each DB Instance. Once an encrypted connection is established, data transferred between the DB Instance and your application will be encrypted during transfer.

NEW QUESTION 283
Which AWS instance address has the following characteristics? :"If you stop an instance, its Elastic IP address is unmapped, and you must remap it when you restart the instance."

A. VPC Addresses
B. EC2 Addresses
C. Both A and B
D. None of the above

Answer: B

NEW QUESTION 285
Please select the most correct answer regarding the persistence of the Amazon Instance Store

A. The data on an instance store volume persists only during the life of the associated Amazon EC2 instance
B. The data on an instance store volume is lost when the security group rule of the associated instance is changed.
C. The data on an instance store volume persists even after associated Amazon EC2 instance is deleted

Answer: A

Explanation:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Storage.html

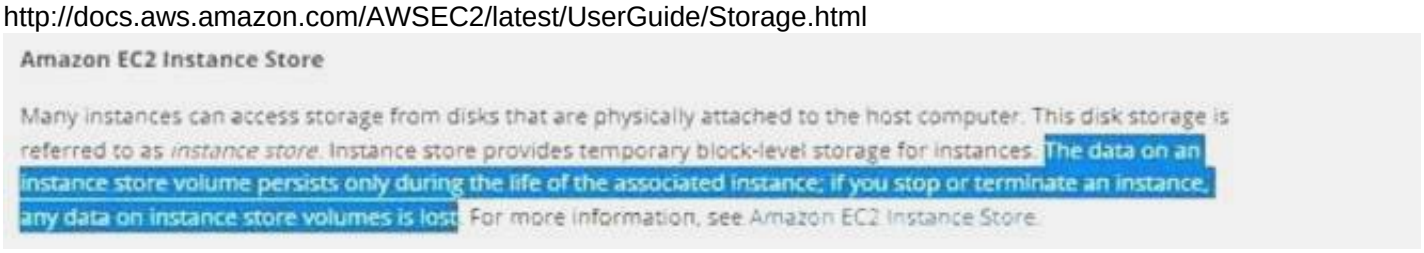
Amazon EC2 Instance Store

Many instances can access storage from disks that are physically attached to the host computer. This disk storage is referred to as *instance store*. Instance store provides temporary block-level storage for instances. The data on an instance store volume persists only during the life of the associated instance; if you stop or terminate an instance, any data on instance store volumes is lost. For more information, see Amazon EC2 Instance Store.

NEW QUESTION 287
Multi-AZ deployment _____ supported for Microsoft SQL Server DB Instances.

A. is not currently
B. is as of 2013
C. is planned to be in 2014
D. will never be

Answer: C

NEW QUESTION 291
Security groups act like a firewall at the instance level, whereas _____ are an additional layer of security that act at the subnet level.

A. DB Security Groups
B. VPC Security Groups
C. network ACLs

Answer: C

NEW QUESTION 296
Is the SQL Server Audit feature supported in the Amazon RDS SQL Server engine?

A. No
B. Yes

Answer: A

Explanation:
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_SQLServer.html

The following server-level permissions are not available on SQL Server DB instances:

- ADMINISTER BULK OPERATIONS
- ALTER ANY CREDENTIAL
- ALTER ANY EVENT NOTIFICATION
- ALTER ANY EVENT SESSION
- ALTER ANY SERVER AUDIT
- ALTER RESOURCES
- ALTER SETTINGS (You can use the DB Parameter Group APIs to modify parameters. For more information, see Working with DB Parameter Groups.
- AUTHENTICATE SERVER
- CONTROL_SERVER
- CREATE DDL EVENT NOTIFICATION
- CREATE ENDPOINT
- CREATE TRACE EVENT NOTIFICATION
- EXTERNAL ACCESS ASSEMBLY
- SHUTDOWN (You can use the RDS reboot option instead)
- UNSAFE ASSEMBLY
- ALTER ANY AVAILABILITY GROUP (SQL Server 2012 only)
- CREATE ANY AVAILABILITY GROUP (SQL Server 2012 only)

NEW QUESTION 301
Which DNS name can only be resolved within Amazon EC2?

A. Internal DNS name
B. External DNS name
C. Global DNS name
D. Private DNS name

Answer: D

Explanation:
To view DNS hostnames for an instance using the console
1. Open the Amazon EC2 console at https://console.aws.amazon.com/ec2/.
2. In the navigation pane, choose Instances.
3. Select your instance from the list.
4. In the details pane, the Public DNS (IPv4) and Private DNS fields display the DNS hostnames, if applicable.

NEW QUESTION 306
If your DB instance runs out of storage space or file system resources, its status will change to _____ and your DB Instance will no longer be available.

A. storage-overflow
B. storage-full
C. storage-exceed
D. storage-overage

Answer: B

Explanation:
https://aws.amazon.com/ko/premiumsupport/knowledge-center/rds-out-of-storage/

Short Description

When an RDS DB instance reaches the STORAGE_FULL state, there is not enough space available for performing basic operations, eventually preventing you from restarting or making connections to the instance.

NEW QUESTION 311
Is it possible to access your EBS snapshots?

A. Yes, through the Amazon S3 APIs.
B. Yes, through the Amazon EC2 APIs.
C. No, EBS snapshots cannot be accessed; they can only be used to create a new EBS volume.
D. EBS doesn't provide snapshot

Answer: B

Explanation:
https://aws.amazon.com/ebs/faqs/?nc1=h_ls
Q: Will I be able to access my snapshots using the regular Amazon S3 API? No, snapshots are only available through the Amazon EC2 API.

NEW QUESTION 315
Does Amazon RDS for SQL Server currently support importing data into the msdb database?

A. No
B. Yes

Answer: A

Explanation:
Amazon RDS for SQL Server does not currently support importing data into the msdb database, though we do support SQL Server Agent jobs. Some SQL Server features that use the msdb database, such as Database Mail and Replication, are not currently supported in Amazon RDS.
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/SQLServer.Procedural.Importing.html#S QLServer.Proced ural.Importing.Procedure

NEW QUESTION 318
Does Route 53 support MX Records?

A. Yes.
B. It supports CNAME records, but not MX records.
C. No
D. Only Primary MX record
E. Secondary MX records are not supported.

Answer: A

Explanation:
http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/ResourceRecordTypes.html#MXForm


MX Format

Each value for an MX resource record set actually contains two values:

- An integer that represents the priority for an email server
- The domain name of the email server

If you specify only one server, the priority can be any integer between 0 and 65535. If you specify multiple servers, the value that you specify for the priority indicates which email server you want email to be routed to first, second, and so on. For example, if you have two email servers and you specify values of 10 and 20 for the priority, email always goes to the server with a priority of 10 unless it's unavailable. If you specify values of 10 and 10, email is routed to the two servers approximately equally.

Example for the Amazon Route 53 console

```
10 mail.example.com
```

Example for the Amazon Route 53 API

```
<Value>10 mail.example.com</Value>
```

NEW QUESTION 322
Select the correct set of steps for exposing the snapshot only to specific AWS accounts

A. Select public for all the accounts and check mark those accounts with whom you want to expose the snapshots and click save.
B. SelectPrivate, enter the IDs of those AWS accounts, and clickSave.
C. SelectPublic, enter the IDs of those AWS accounts, and clickSave.
D. SelectPublic, mark the IDs of those AWS accounts as private, and clickSav

Answer: B

Explanation:
“To expose the snapshot to only specific AWS accounts, choose Private, enter the ID of the AWS account (without hyphens) in the AWS Account Number field, and choose Add Permission. Repeat until you’ve added all the required AWS accounts” http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-modifying-snapshot-permissions.html

NEW QUESTION 325
In the context of MySQL, version numbers are organized as MySQL version = X.Y.Z. What does X denote here?

A. release level
B. minor version
C. version number
D. major version

Answer: D

Explanation:

MySQL on Amazon RDS Versions

For MySQL, version numbers are organized as version = X.Y.Z. In Amazon RDS terminology, X.Y denotes the major version, and Z is the minor version number. For Amazon RDS implementations, a version change is considered major if the major version number changes—for example, going from version 5.6 to 5.7. A version change is considered minor if only the minor version number changes—for example, going from version 5.6.22 to 5.6.23.

Amazon RDS currently supports MySQL major versions 5.5, 5.6, and 5.7. MySQL minor version support varies by AWS Region. Use the following table to see what MySQL minor versions are supported in each AWS Region.

NEW QUESTION 327
It is advised that you watch the Amazon CloudWatch " _____ " metric (available via the AWS Management Console or Amazon Cloud Watch APIs) carefully and recreate the Read Replica should it fall behind due to replication errors.

A. Write Lag
B. Read Replica
C. Replica Lag
D. Single Replica

Answer: C

Explanation:
The amount of time a Read Replica DB instance lags behind the source DB instance. Applies to MySQL, MariaDB, and PostgreSQL Read Replicas.
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/rds-metricscollected.html

ReplicaLag	The amount of time a Read Replica DB instance lags behind the source DB instance. Applies to MySQL, MariaDB, and PostgreSQL Read Replicas. Units: Seconds

NEW QUESTION 328
By default, what are ENIs that are automatically created and attached to instances using the EC2 console set to do when the attached instance terminates?

A. Remain as is
B. Terminate
C. Hibernate
D. Pause

Answer: B

Explanation:
By default, elastic network interfaces that are automatically created and attached to instances using the console are set to terminate when the instance terminates. However, network interfaces created using the command line interface aren't set to terminate when the instance terminates.

NEW QUESTION 330
Are you able to integrate a multi-factor token service with the AWS Platform?

A. Yes, you can integrate private multi-factor token devices to authenticate users to the AWS platform.
B. No, you cannot integrate multi-factor token devices with the AWS platform.
C. Yes, using the AWS multi-factor token devices to authenticate users on the AWS platfor

Answer: C

Explanation:
Private MFA does not apply here.
Q. What is AWS MFA?
AWS multi-factor authentication (AWS MFA) provides an extra level of security that you can apply to your AWS environment. You can enable AWS MFA for your AWS account and for individual AWS Identity and Access Management (IAM) users you create under your account.

NEW QUESTION 335
You can use _____ and _____ to help secure the instances in your VPC,

A. security groups and multi-factor authentication
B. security groups and 2-Factor authentication
C. security groups and biometric authentication
D. security groups and network ACLs

Answer: D

NEW QUESTION 338
Fill in the blanks: is a durable, block-level storage volume that you can attach to a single, running Amazon EC2 instance.

A. Amazon S3
B. Amazon EBS
C. None of these
D. All of these

Answer: B

NEW QUESTION 339
If I want my instance to run on a single-tenant hardware, which value do I have to set the instance's tenancy attribute to?

A. dedicated
B. isolated
C. one
D. reserved

Answer: A

Explanation:
http://aws.amazon.com/ec2/dedicated-hosts/

NEW QUESTION 341
What does Amazon ELB stand for?

A. Elastic Linux Box.
B. Encrypted Linux Box.
C. Encrypted Load Balancing.
D. Elastic Load Balancin

Answer: D

NEW QUESTION 343
Is there a limit to the number of groups you can have?

A. Yes for all users except root
B. No
C. Yes, unless special permission granted
D. Yes for all users

Answer: D

Explanation:
Currently you can request to increase the limit on users per AWS account, groups per AWS account, roles per AWS account, instance profiles per AWS account, and server certificates per
AWS account. http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_iam-limits.html

NEW QUESTION 347
Is there any way to own a direct connection to Amazon Web Services?

A. You can create an encrypted tunnel to VPC, but you don't own the connection.
B. Yes, it's called Amazon Dedicated Connection.
C. No, AWS only allows access from the public Internet.
D. Yes, it's called Direct Connec

Answer: D

NEW QUESTION 350
What is the maximum response time for a Business level Premium Support case?

A. 30 minutes
B. 1 hour
C. 12 hours
D. 10 minutes

Answer: B

NEW QUESTION 354
Does Dynamic DB support in-place atomic updates?

A. It is not defined
B. No

C. Yes
D. It does support in-place non-atomic updates

Answer: C

Explanation:
Q: Does DynamoDB support in-place atomic updates?
Amazon DynamoDB supports fast in-place updates. You can increment or decrement a numeric attribute in a row using a single API call. Similarly, you can atomically add or remove to sets, lists, or maps.
https://aws.amazon.com/dynamodb/faqs/

NEW QUESTION 359
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your SAA-C01 Exam with Our Prep Materials Via below:

https://www.certleader.com/SAA-C01-dumps.html