Exam Questions SAA-C01

AWS Certified Solutions Architect - Associate

https://www.2passeasy.com/dumps/SAA-C01/

NEW QUESTION 1
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 2
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 3
You have launched an EC2 instance with four (4) 500 GB EBS Provisioned IOPS volumes attached The EC2 Instance Is EBS-Optimized and supports 500 Mbps throughput between EC2 and EBS. The two EBS volumes are configured as a single RAID o device, and each Provisioned IOPS volume is provisioned with 4.000 IOPS (4 000 16KB reads or writes) for a total of 16.000 random IOPS on the instance The EC2 Instance initially delivers the expected 16 000 IOPS random read and write performance Sometime later in order to increase the total random I/O performance of the instance,
you add an additional two 500 GB EBS Provisioned IOPS volumes to the RAID Each volume Is provisioned to 4.000 IOPs like the original four for a total of 24.000 IOPS on the EC2 instance Monitoring shows that the EC2 instance CPU utilization increased from 50% to 70%. but the total random IOPS measured at the instance level does not increase at all.
What is the problem and a valid solution?
A. Larger storage volumes support higher Provisioned IOPS rates: increase the provisioned volume storage of each of the 6 EBS volumes to 1TB
B. The EBS-Optimized throughput limits the total IOPS that can be utilized use an EBS-Optimized instance that provides larger throughput.
C. Small block sizes cause performance degradation, limiting the I'O throughput, configure the instance device driver and file system to use 64KB blocks to increase throughput.
D. RAID 0 only scales linearly to about 4 devices, use RAID 0 with 4 EBS Provisioned IOPS volumes but increase each Provisioned IOPS EBS volume to 6.000 IOPS.
E. The standard EBS instance root volume limits the total IOPS rate, change the instant root volume to also be a 500GB 4.000 Provisioned IOPS volume.

Answer: E

NEW QUESTION 4
Your company is in the process of developing a next generation pet collar that collects biometric information to assist families with promoting healthy lifestyles for their pets Each collar will push 30kb of biometric data In JSON format every 2 seconds to a collection platform that will process and analyze the data providing health trending information back to the pet owners and veterinarians via
a web portal Management has tasked you to architect the collection platform ensuring the following requirements are met.
Provide the ability for real-time analytics of the inbound biometric data Ensure processing of the biometric data is highly durable. Elastic and parallel The results of the analytic processing should be persisted for data mining
Which architecture outlined below win meet the initial requirements for the collection platform?

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

Answer: B

Explanation:
The POC solution is being scaled up by 1000, which means it will require 72TB of Storage to retain 24 months' worth of data. This rules out RDS as a possible DB solution which leaves you with RedShift. I believe DynamoDB is a more cost effective and scales better for ingest rather than using EC2 in an auto scaling group. Also, this example solution from AWS is somewhat similar for reference.
http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_timeseriesprocessing_16.pdf

NEW QUESTION 5
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 6
Amazon EC2 provides virtual computing environments known as _____ .

A. instances
B. volumes
C. microsystems
D. servers

Answer: A

Explanation:
Amazon EC2 provides virtual computing environments known as instances. When you launch an instance, the instance type that you specify determines the hardware of the host computer used for your instance. Each instance type offers different compute, memory, and storage capabilities and are grouped in instance families based on these capabilities. Select an instance type based on the requirements of the application or software that you plan to run on your instance.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts.html

NEW QUESTION 7
Your company runs a customer facing event registration site This site is built with a 3-tier architecture with web and application tier servers and a MySQL database The application requires 6 web tier servers and 6 application tier servers for normal operation, but can run on a minimum of 65% server capacity and a single MySQL database. When deploying this application in a region with three availability zones (AZs) which architecture provides high availability?

A. A web tier deployed across 2 AZs with 3 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer), and an application tier deployed across 2 AZs with 3 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB, and one RDS (Relational Database Service) instance deployed with read replicas in the other AZ.
B. A web tier deployed across 3 AZs with 2 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer) and an application tier deployed across 3 AZs with 2 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and one RDS (Relational Database Service) Instance deployed with read replicas in the two other AZs.
C. A web tier deployed across 2 AZs with 3 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer) and an application tier deployed across 2 AZs with 3 EC2 instances m each AZ inside an Auto Scaling Group behind an ELS and a Multi-AZ RDS (Relational Database Service) deployment.
D. A web tier deployed across 3 AZs with 2 EC2 (Elastic Compute Cloud) instances in each AZ Inside an Auto Scaling Group behind an ELB (elastic load balancer). And an application tier deployed across 3 AZs with 2 EC2 instances in each AZ inside an Auto Scaling Group behind an EL
E. And a Multi-AZ RDS (Relational Database services) deployment.

Answer: D

Explanation:
Amazon RDS Multi-AZ Deployments
Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a Multi- AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure (for example, instance hardware failure, storage failure, or network disruption), Amazon RDS performs an automatic failover to the standby, so that you can resume database operations as soon as the failover is complete. Since the endpoint for your DB Instance remains the same after a failover, your application can resume database operation without the need for manual administrative intervention.
Enhanced Durability
Multi-AZ deployments for the MySQL, Oracle, and PostgreSQL engines utilize synchronous physical replication to keep data on the standby up-to-date with the primary. Multi-AZ deployments for the SQL Server engine use synchronous logical replication to achieve the same result, employing SQL Server-native Mirroring technology. Both approaches safeguard your data in the event of a DB Instance failure or loss of an Availability Zone.
If a storage volume on your primary fails in a Multi-AZ deployment, Amazon RDS automatically initiates a failover to the up-to-date standby. Compare this to a Single-AZ deployment: in case of a Single-AZ database failure, a user-initiated point-in-time-restore operation will be required. This operation can take several hours to complete, and any data updates that occurred after the latest restorable time (typically within the last five minutes) will not be available.
Amazon Aurora employs a highly durable, SSD-backed virtualized storage layer purpose-built for database workloads. Amazon Aurora automatically replicates

your volume six ways, across three Availability Zones. Amazon Aurora storage is fault-tolerant, transparently handling the loss of up to two copies of data without affecting database write availability and up to three copies without affecting read availability. Amazon Aurora storage is also self-healing. Data blocks and disks are continuously scanned for errors and replaced automatically.
Increased Availability
You also benefit from enhanced database availability when running Multi-AZ deployments. If an Availability Zone failure or DB Instance failure occurs, your availability impact is limited to the time automatic failover takes to complete: typically under one minute for Amazon Aurora and one to two minutes for other database engines (see the RDS FAQ for details).
The availability benefits of Multi-AZ deployments also extend to planned maintenance and backups. In the case of system upgrades like OS patching or DB Instance scaling, these operations are applied first on the standby, prior to the automatic failover. As a result, your availability impact is, again, only the time required for automatic failover to complete.
Unlike Single-AZ deployments, I/O activity is not suspended on your primary during backup for Multi- AZ deployments for the MySQL, Oracle, and PostgreSQL engines, because the backup is taken from the standby. However, note that you may still experience elevated latencies for a few minutes during backups for Multi-AZ deployments.
On instance failure in Amazon Aurora deployments, Amazon RDS uses RDS Multi-AZ technology to automate failover to one of up to 15 Amazon Aurora Replicas you have created in any of three Availability Zones. If no Amazon Aurora Replicas have been provisioned, in the case of a failure, Amazon RDS will attempt to create a new Amazon Aurora DB instance for you automatically. https://www.airpair.com/aws/posts/building-a-scalable-web-app-on-amazon-web-services-p1

NEW QUESTION 8
Your application is using an ELB in front of an Auto Scaling group of web/application servers deployed across two AZs and a Multi-AZ RDS Instance for data persistence.
The database CPU is often above 80% usage and 90% of I/O operations on the database are reads. To improve performance you recently added a single-node Memcached ElastiCache Cluster to cache frequent DB query results. In the next weeks the overall workload is expected to grow by 30%.
Do you need to change anything in the architecture to maintain the high availability or the application with the anticipated additional load? Why?

A. Yes, you should deploy two Memcached ElastiCache Clusters in different AZs because the RDS instance will not be able to handle the load if the cache node fails.
B. No, if the cache node fails you can always get the same data from the DB withouthaving any availability impact.
C. No, if the cache node fails the automated ElastiCache node recovery feature will prevent any availability impact.
D. Yes, you should deploy the Memcached ElastiCache Cluster with two nodes in the same AZ as the RDS DB master instance to handle the load if one cache node fails.

Answer: A

Explanation:
A single-node Memcached ElastiCache cluster failure is nothing but a total failure. (Even though AWS will automatically recover the failed node, there are no other nodes in the cluster) http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/BestPractices.html Mitigating Node Failures
To mitigate the impact of a node failure, spread your cached data over more nodes. Because Memcached does not support replication, a node failure will always result in some data loss from your cluster.
When you create your Memcached cluster you can create it with 1 to 20 nodes, or more by special request. Partitioning your data across a greater number of nodes means you'll lose less data if a node fails. For example, if you partition your data across 10 nodes, any single node stores approximately 10% of your cached data. In this case, a node failure loses approximately 10% of your cache which needs to be replaced when a replacement node is created and provisioned.
Mitigating Availability Zone Failures
To mitigate the impact of an availability zone failure, locate your nodes in as many availability zones as possible. In the unlikely event of an AZ failure, you will lose only the data cached in that AZ, not the data cached in the other AZs.

NEW QUESTION 9
An International company has deployed a multi-tier web application that relies on DynamoDB in a single region. For regulatory reasons they need disaster recovery capability In a separate region with a Recovery Time Objective of 2 hours and a Recovery Point Objective of 24 hours They should synchronize their data on a regular basis and be able to provision me web application rapidly using CloudFormation.
The objective is to minimize changes to the existing web application, control the throughput of DynamoDB used for the synchronization of data and synchronize only the modified elements. Which design would you choose to meet these requirements?

A. Use AWS data Pipeline to schedule a DynamoDB cross region copy once a da
B. create a Lastupdated' attribute in your DynamoDB table that would represent the timestamp of the last update and use it as a filter.
C. Use EMR and write a custom script to retrieve data from DynamoDB in the current region using a SCAN operation and push it to DynamoDB in the second region.
D. Use AWS data Pipeline to schedule an export of the DynamoDB table to S3 in the current region once a day then schedule another task immediately after it that will import data from S3 to DynamoDB in the other region.
E. Send also each Ante into an SQS queue in me second region; use an auto-scaling group behind the SQS queue to replay the write in the second region.

Answer: A

NEW QUESTION 10
An ERP application is deployed across multiple AZs in a single region. In the event of failure, the Recovery Time Objective (RTO) must be less than 3 hours, and the Recovery Point Objective (RPO) must be 15 minutes the customer realizes that data corruption occurred roughly 1.5 hours ago.
What DR strategy could be used to achieve this RTO and RPO in the event of this kind of failure?

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

Answer: A

NEW QUESTION 10
Your startup wants to implement an order fulfillment process for selling a personalized gadget that needs an average of 3-4 days to produce with some orders taking up to 6 months you expect 10 orders per day on your first day. 1000 orders per day after 6 months and 10,000 orders after 12 months.
Orders coming in are checked for consistency men dispatched to your manufacturing plant for production quality control packaging shipment and payment processing If the product does not meet the quality standards at any stage of the process employees may force the process to repeat a step Customers are

notified via email about order status and any critical issues with their orders such as payment failure.
Your case architecture includes AWS Elastic Beanstalk for your website with an RDS MySQL instance for customer data and orders.
How can you implement the order fulfillment process while making sure that the emails are delivered reliably?

A. Add a business process management application to your Elastic Beanstalk app servers and re-use the ROS database for tracking order status use one of the Elastic Beanstalk instances to send emails to customers.
B. Use SWF with an Auto Scaling group of activity workers and a decider instance in another Auto Scaling group with min/max=1 Use the decider instance to send emails to customers.
C. Use SWF with an Auto Scaling group of activity workers and a decider instance in another Auto Scaling group with min/max=1 use SES to send emails to customers.
D. Use an SQS queue to manage all process tasks Use an Auto Scaling group of EC2 Instances that poll the tasks and execute the
E. Use SES to send emails to customers.

Answer: C

Explanation:
http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_ecommerce_checkout_13.pdf

NEW QUESTION 13
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 15
Your system recently experienced down time during the troubleshooting process. You found that a new administrator mistakenly terminated several production EC2 instances.
Which of the following strategies will help prevent a similar situation in the future? The administrator still must be able to:
- launch, start stop, and terminate development resources.
- launch and start production instances.

A. Create an IAM user, which is not allowed to terminate instances by leveraging production EC2 termination protection.
B. Leverage resource based tagging along with an IAM user, which can prevent specific users from terminating production EC2 resources.
C. Leverage EC2 termination protection and multi-factor authentication, which together require users to authenticate before terminating EC2 instances
D. Create an IAM user and apply an IAM role which prevents users from terminating production EC2 instances.

Answer: B

Explanation:
Working with volumes
When an API action requires a caller to specify multiple resources, you must create a policy statement that allows users to access all required resources. If you need to use a Condition element with one or more of these resources, you must create multiple statements as shown in this example. The following policy allows users to attach volumes with the tag "volume_user=iam-user-name" to
instances with the tag "department=dev", and to detach those volumes from those instances. If you attach this policy to an IAM group, the aws:username policy variable gives each IAM user in the group permission to attach or detach volumes from the instances with a tag named volume_user that has his or her IAM user name as a value.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "ec2:AttachVolume",
      "ec2:DetachVolume"
    ],
    "Resource": "arn:aws:ec2:us-east-1:123456789012:instance/*",
    "Condition": {
      "StringEquals": {
        "ec2:ResourceTag/department": "dev"
      }
    }
  },
  {
    "Effect": "Allow",
    "Action": [
      "ec2:AttachVolume",
      "ec2:DetachVolume"
    ],
    "Resource": "arn:aws:ec2:us-east-1:123456789012:volume/*",
    "Condition": {
      "StringEquals": {
        "ec2:ResourceTag/volume_user": "${aws:username}"
      }
    }
  }
 ]
}
```

Launching instances (RunInstances)

The RunInstances API action launches one or more instances. RunInstances requires an AMI and creates an instance; and users can specify a key pair and security group in the request. Launching into EC2-VPC requires a subnet, and creates a network interface. Launching from an Amazon EBS-backed AMI creates a volume. Therefore, the user must have permission to use these Amazon EC2
resources. The caller can also configure the instance using optional parameters to RunInstances, such as the instance type and a subnet. You can create a policy statement that requires users to specify an optional parameter, or restricts users to particular values for a parameter. The examples in this section demonstrate some of the many possible ways that you can control the configuration of an instance that a user can launch.

Note that by default, users don't have permission to describe, start, stop, or terminate the resulting instances. One way to grant the users permission to manage the resulting instances is to create a specific tag for each instance, and then create a statement that enables them to manage instances with that tag. For more information, see 2: Working with instances.

/a. AMI

The following policy allows users to launch instances using only the AMIs that have the specified tag, "department=dev", associated with them. The users can't launch instances using other AMIs because the Condition element of the first statement requires that users specify an AMI that has this tag. The users also can't launch into a subnet, as the policy does not grant permissions for the subnet and network interface resources. They can, however, launch into EC2-Classic. The second statement uses a wildcard to enable users to create instance resources, and requires users to specify the key pair project_keypair and the security group sg-1a2b3c4d. Users are still able to launch instances without a key pair.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": "ec2:RunInstances",
    "Resource": [
      "arn:aws:ec2:region::image/ami-*"
    ],
    "Condition": {
      "StringEquals": {
        "ec2:ResourceTag/department": "dev"
      }
    }
  },
  {
    "Effect": "Allow",
    "Action": "ec2:RunInstances",
    "Resource": [
      "arn:aws:ec2:region:account:instance/*",
      "arn:aws:ec2:region:account:volume/*",
      "arn:aws:ec2:region:account:key-pair/project_keypair",
      "arn:aws:ec2:region:account:security-group/sg-1a2b3c4d"
    ]
  }
 ]
}
```

Alternatively, the following policy allows users to launch instances using only the specified AMIs, ami-9e1670f7 and ami-45cf5c3c. The users can't launch an instance using other AMIs (unless another statement grants the users permission to do so), and the users can't launch an instance into a subnet.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": "ec2:RunInstances",
    "Resource": [
      "arn:aws:ec2:region::image/ami-9e1670f7",
      "arn:aws:ec2:region::image/ami-45cf5c3c",
      "arn:aws:ec2:region:account:instance/*",
      "arn:aws:ec2:region:account:volume/*",
      "arn:aws:ec2:region:account:key-pair/*",
      "arn:aws:ec2:region:account:security-group/*"
    ]
   }
  ]
}
```

Alternatively, the following policy allows users to launch instances from all AMIs owned by Amazon. The Condition element of the first statement tests whether ec2:Owner is amazon. The users can't launch an instance using other AMIs (unless another statement grants the users permission to do so). The users are able to launch an instance into a subnet.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": "ec2:RunInstances",
    "Resource": [
      "arn:aws:ec2:region::image/ami-*"
    ],
    "Condition": {
      "StringEquals": {
        "ec2:Owner": "amazon"
        }
    }
  },
  {
    "Effect": "Allow",
    "Action": "ec2:RunInstances",
    "Resource": [
      "arn:aws:ec2:region:account:instance/*",
      "arn:aws:ec2:region:account:subnet/*",
      "arn:aws:ec2:region:account:volume/*",
      "arn:aws:ec2:region:account:network-interface/*",
      "arn:aws:ec2:region:account:key-pair/*",
      "arn:aws:ec2:region:account:security-group/*"
      ]
    }
  ]
}
```

/b. Instance type

The following policy allows users to launch instances using only the t2.micro or t2.small instance type, which you might do to control costs. The users can't launch larger instances because the Condition element of the first statement tests whether ec2:InstanceType is either t2.micro or t2.small.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": "ec2:RunInstances",
    "Resource": [
      "arn:aws:ec2:region:account:instance/*"
    ],
    "Condition": {
      "StringEquals": {
        "ec2:InstanceType": ["t2.micro", "t2.small"]
      }
    }
  },
  {
    "Effect": "Allow",
    "Action": "ec2:RunInstances",
    "Resource": [
      "arn:aws:ec2:region::image/ami-*",
      "arn:aws:ec2:region:account:subnet/*",
      "arn:aws:ec2:region:account:network-interface/*",
      "arn:aws:ec2:region:account:volume/*",
      "arn:aws:ec2:region:account:key-pair/*",
      "arn:aws:ec2:region:account:security-group/*"
      ]
    }
  ]
}
```

Alternatively, you can create a policy that denies users permission to launch any instances except t2.micro and t2.small instance types.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Deny",
    "Action": "ec2:RunInstances",
    "Resource": [
      "arn:aws:ec2:region:account:instance/*"
    ],
    "Condition": {
      "StringNotEquals": {
        "ec2:InstanceType": ["t2.micro", "t2.small"]
        }
    }
  },
  {
    "Effect": "Allow",
    "Action": "ec2:RunInstances",
    "Resource": [
      "arn:aws:ec2:region::image/ami-*",
      "arn:aws:ec2:region:account:network-interface/*",
      "arn:aws:ec2:region:account:instance/*",
      "arn:aws:ec2:region:account:subnet/*",
      "arn:aws:ec2:region:account:volume/*",
      "arn:aws:ec2:region:account:key-pair/*",
      "arn:aws:ec2:region:account:security-group/*"
      ]
    }
  ]
}
```

/c. Subnet

The following policy allows users to launch instances using only the specified subnet, subnet- 12345678. The group can't launch instances into any another subnet (unless another statement grants the users permission to do so). Users are still able to launch instances into EC2-Classic.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": "ec2:RunInstances",
    "Resource": [
      "arn:aws:ec2:region:account:subnet/subnet-12345678",
      "arn:aws:ec2:region:account:network-interface/*",
      "arn:aws:ec2:region:account:instance/*",
      "arn:aws:ec2:region:account:volume/*",
      "arn:aws:ec2:region::image/ami-*",
      "arn:aws:ec2:region:account:key-pair/*",
      "arn:aws:ec2:region:account:security-group/*"
    ]
  }
  ]
}
```

Alternatively, you could create a policy that denies users permission to launch an instance into any other subnet. The statement does this by denying permission to create a network interface, except where subnet subnet-12345678 is specified. This denial overrides any other policies that are created to allow launching instances into other subnets. Users are still able to launch instances into EC2- Classic.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Deny",
    "Action": "ec2:RunInstances",
    "Resource": [
      "arn:aws:ec2:region:account:network-interface/*"
    ],
    "Condition": {
      "ArnNotEquals": {
        "ec2:Subnet": "arn:aws:ec2:region:account:subnet/subnet-12345678
        }
    }
  },
  {
    "Effect": "Allow",
    "Action": "ec2:RunInstances",
    "Resource": [
      "arn:aws:ec2:region::image/ami-*",
      "arn:aws:ec2:region:account:network-interface/*",
      "arn:aws:ec2:region:account:instance/*",
      "arn:aws:ec2:region:account:subnet/*",
      "arn:aws:ec2:region:account:volume/*",
      "arn:aws:ec2:region:account:key-pair/*",
      "arn:aws:ec2:region:account:security-group/*"
      ]
    }
  ]
}
```

https://aws.amazon.com/blogs/security/resource-level-permissions-for-ec2-controllingmanagement- access-on-specific-instances/
August 2016 Update One way to work around this is to use a combination of an Amazon CloudWatch Events rule and AWS Lambda to tag newly created instances.

NEW QUESTION 16
Your company previously configured a heavily used, dynamically routed VPN connection between your on-premises data center and AWS. You recently provisioned a DirectConnect connection and would like to start using the new connection. After configuring DirectConnect settings in the AWS Console, which of the following options win provide the most seamless transition for your users?

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

Answer: C

Explanation:
Q. Can I use AWS Direct Connect and a VPN Connection to the same VPC simultaneously?
Yes. However, only in fail-over scenarios. The Direct Connect path will always be preferred, when established, regardless of AS path prepending.
https://aws.amazon.com/directconnect/faqs/

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 24
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 25
You are migrating a legacy client-server application to AWS. The application responds to a specific DNS domain (e.g. www.example.com) and has a 2-tier architecture, with multiple application servers and a database server. Remote clients use TCP to connect to the application servers. The application servers need to know the IP address of the clients in order to function properly and are currently taking that information from the TCP socket. A Multi-AZ RDS MySQL instance will be used for the database.
During the migration you can change the application code, but you have to file a change request. How would you implement the architecture on AWS in order to maximize scalability and high availability?

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

Answer: D

NEW QUESTION 29
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 31
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 32
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 36
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 40
You currently operate a web application. In the AWS US-East region The application runs on an autoscaled layer of EC2 instances and an RDS Multi-AZ database Your IT security compliance officer has
tasked you to develop a reliable and durable logging solution to track changes made to your EC2.IAM And RDS resources. The solution must ensure the integrity and confidentiality of your log data. Which of these solutions would you recommend?

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

Answer: A

NEW QUESTION 44
Your company has recently extended its datacenter into a VPC on AVVS to add burst computing capacity as needed Members of your Network Operations Center need to be able to go to the AWS Management Console and administer Amazon EC2 instances as necessary You don't want to create new IAM users for each NOC member and make those users sign in again to the AWS Management Console Which option below will meet the needs for your NOC members?

A. Use OAuth 2 0 to retrieve temporary AWS security credentials to enable your NOC members to sign in to the AVVS Management Console.
B. Use web Identity Federation to retrieve AWS temporary security credentials to enable your NOC members to sign in to the AWS Management Console.
C. Use your on-premises SAML 2.0-compliant identity provider (IDP) to grant the NOC members federated access to the AWS Management Console via the AWS single sign-on (SSO) endpoint.
D. Use your on-premises SAML2.0-compliant identity provider (IDP) to retrieve temporary security credentials to enable NOC members to sign in to the AWS Management Console.

Answer: C

NEW QUESTION 46
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 49
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 53
An AWS customer runs a public blogging website. The site users upload two million blog entries a month. The average blog entry size is 200 KB. The access rate to blog entries drops to negligible 6 months after publication and users rarely access a blog entry 1 year after publication. Additionally,
blog entries have a high update rate during the first 3 months following publication, this drops to no updates after 6 months. The customer wants to use CloudFront to improve his user's load times. Which of the following recommendations would you make to the customer?

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

Answer: C

NEW QUESTION 55
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 56
A company is running a batch analysis every hour on their main transactional DB. running on an RDS MySQL instance to populate their central Data Warehouse running on Redshift During the execution of the batch their transactional applications are very slow When the batch completes they need to update the top management dashboard with the new data The dashboard is produced by another system running on-premises that is currently started when a manually-sent email notifies that an update is required The on-premises system cannot be modified because is managed by another team.
How would you optimize this scenario to solve performance issues and automate the process as much as possible?

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

Answer: C

Explanation:
If you want to prevent your reporting and analytic processing from interfering with the performance of your OLTP workload."
If I understand the above statement correctly, they are saying to separate reporting and analytic processing from OLTP. In other word, use RedShift for reporting and analytic processing and use RDS for OLTP workload.

NEW QUESTION 61
You are running a news website in the eu-west-1 region that updates every 15 minutes. The website has a world-wide audience it uses an Auto Scaling group behind an Elastic Load Balancer and an Amazon RDS database Static content resides on Amazon S3, and is distributed through Amazon CloudFront. Your Auto Scaling group is set to trigger a scale up event at 60% CPU utilization, you use an Amazon RDS extra large DB instance with 10.000 Provisioned IOPS its CPU utilization is around 80%. While freeable memory is in the 2 GB range.
Web analytics reports show that the average load time of your web pages is around 1 5 to 2 seconds,
but your SEO consultant wants to bring down the average load time to under 0.5 seconds. How would you improve page load times for your users? (Choose three.)

A. Lower the scale up trigger of your Auto Scaling group to 30% so it scales more aggressively.
B. Add an Amazon ElastiCache caching layer to your application for storing sessions and frequent DB queries
C. Configure Amazon CloudFront dynamic content support to enable caching of re-usable content from your site
D. Switch Amazon RDS database to the high memory extra large Instance type
E. Set up a second installation in another region, and use the Amazon Route 53 latency-based routing feature to select the right region.

Answer: BCE

Explanation:
The freeable memory includes the amount of physical memory left unused by the system plus the total amount of buffer or page cache memory that are free and available.
So it's freeable memory across the entire system. While MySQL is the main consumer of memory on the host we do have internal processes in addition to the OS that use up a small amount of additional memory.
If you see your freeable memory near 0 or also start seeing swap usage then you may need to scale up to a larger instance class or adjust MySQL memory settings. For example decreasing
the innodb_buffer_pool_size (by default set to 75% of physical memory) is one way example of adjusting MySQL memory settings
Takeaway: extra mem is not going to help page load times here, but a 2nd region might. Keep in mind they're going for a 66%-75% reduction in page load times – what if you added a region in Australia or HK, would that not help your worldwide users? rather than having traffic go to us-east.

NEW QUESTION 63
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 66
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 68
Your team has a tomcat-based Java application you need to deploy into development, test and production environments. After some research, you opt to use Elastic Beanstalk due to its tight integration with your developer tools and RDS due to its ease of management. Your QA team lead points out that you need to roll a sanitized set of production data into your environment on a nightly basis. Similarly, other software teams in your org want access to that same restored data via their EC2 instances in your VPC. The optimal setup for persistence and security that meets the above requirements would be the following.

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

Answer: C

Explanation:
Elastic Beanstalk provides support for running Amazon RDS instances in your Elastic Beanstalk environment. This works great for development and testing environments, but is not ideal for a
production environment because it ties the lifecycle of the database instance to the lifecycle of your application's environment.
It can't be D because RDS is opened to all "hosts in your application subnets" where C only opens RDS to specific client machines in a specific security group.
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/AWSHowTo.RDS.html

NEW QUESTION 70
Your customer is willing to consolidate their log streams (access logs application logs security logs etc.) in one single system. Once consolidated, the customer wants to analyze these logs in real time based on heuristics. From time to time, the customer needs to validate heuristics, which requires going back to data samples extracted from the last 12 hours?
What is the best approach to meet your customer's requirements?
A. Send all the log events to Amazon SQ
B. Setup an Auto Scaling group of EC2 servers to consume the logs and apply the heuristics.
C. Send all the log events to Amazon Kinesis develop a client process to apply heuristics on the logs
D. Configure Amazon Cloud Trail to receive custom logs, use EMR to apply heuristics the logs
E. Setup an Auto Scaling group of EC2 syslogd servers, store the logs on S3 use EMR to apply heuristics on the logs

Answer: B

Explanation:
Amazon Kinesis Streams allows for real-time data processing. With Amazon Kinesis Streams, you can continuously collect data as it is generated and promptly react to critical information about your business and operations.
https://aws.amazon.com/kinesis/streams/

NEW QUESTION 75
You deployed your company website using Elastic Beanstalk and you enabled log file rotation to S3. An Elastic Map Reduce job is periodically analyzing the logs on S3 to build a usage dashboard that you share with your CIO.
You recently improved overall performance of the website using Cloud Front for dynamic content delivery and your website as the origin.
After this architectural change, the usage dashboard shows that the traffic on your website dropped by an order of magnitude. How do you fix your usage dashboard'?

A. Enable Cloud Front to deliver access logs to S3 and use them as input of the Elastic Map Reduce job.
B. Turn on Cloud Trail and use trail log tiles on S3 as input of the Elastic Map Reduce job
C. Change your log collection process to use Cloud Watch ELB metrics as input of the Elastic Map Reduce job
D. Use Elastic Beanstalk "Rebuild Environment" option to update log delivery to the Elastic Map Reduce job.
E. Use Elastic Beanstalk 'Restart App server(s)" option to update log delivery to the Elastic Map Reduce job.

Answer: A

NEW QUESTION 79
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 81
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 85
Select the most correct answer
The device name /dev/sda1 (within Amazon EC2) is _____

A. Possible for EBS volumes
B. Reserved for the root device
C. Recommended for EBS volumes
D. Recommended for instance store volumes

Answer: B

Explanation:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/RootDeviceStorage.html The root device is typically /dev/sda1 (Linux) or xvda (Windows).

NEW QUESTION 88
What is the Reduced Redundancy option in Amazon S3?

A. Less redundancy for a lower cost.
B. It doesn't exist in Amazon S3, but in Amazon EBS.
C. It allows you to destroy any copy of your files outside a specific jurisdiction.
D. It doesn't exist at all

Answer: A

NEW QUESTION 89
While creating an Amazon RDS DB, your first task is to set up a DB that controls what IP addresses or EC2 instances have access to your DB Instance.

A. Security Pool
B. Secure Zone
C. Security Token Pool
D. Security Group

Answer: D

NEW QUESTION 90
Every user you create in the IAM system starts with .

A. Partial permissions
B. Full permissions
C. No permissions

Answer: C

NEW QUESTION 93
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 94
Amazon SWF is designed to help users...

A. Design graphical user interface interactions
B. Manage user identification and authorization
C. Store Web content
D. Coordinate synchronous and asynchronous tasks which are distributed and fault toleran

Answer: D

NEW QUESTION 99
By default, EBS volumes that are created and attached to an instance at launch are deleted when that instance is terminated. You can modify this behavior by changing the value of the flag to false when you launch the instance

A. DeleteOnTermination
B. RemoveOnDeletion
C. RemoveOnTermination
D. TerminateOnDeletion

Answer: A

Explanation:
By default, Amazon EBS root device volumes are automatically deleted when the instance terminates. However, by default, any additional EBS volumes that you attach at launch, or any EBS volumes that you attach to an existing instance persist even after the instance terminates.
This behavior is controlled by the volume's DeleteOnTermination attribute, which you can modify.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/terminating-instances.html

NEW QUESTION 100
Will my standby RDS instance be in the same Region as my primary?

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

Answer: B

Explanation:

Q: Will my standby be in the same Region as my primary?

Yes. Your standby is automatically provisioned in a different Availability Zone of the *same Region* as your DB instance primary.

NEW QUESTION 102
True or False: When using IAM to control access to your RDS resources, the key names that can be used are case sensitive. For example, aws:CurrentTime is NOT equivalent to AWS:currenttime.

A. TRUE
B. FALSE

Answer: A

Explanation:

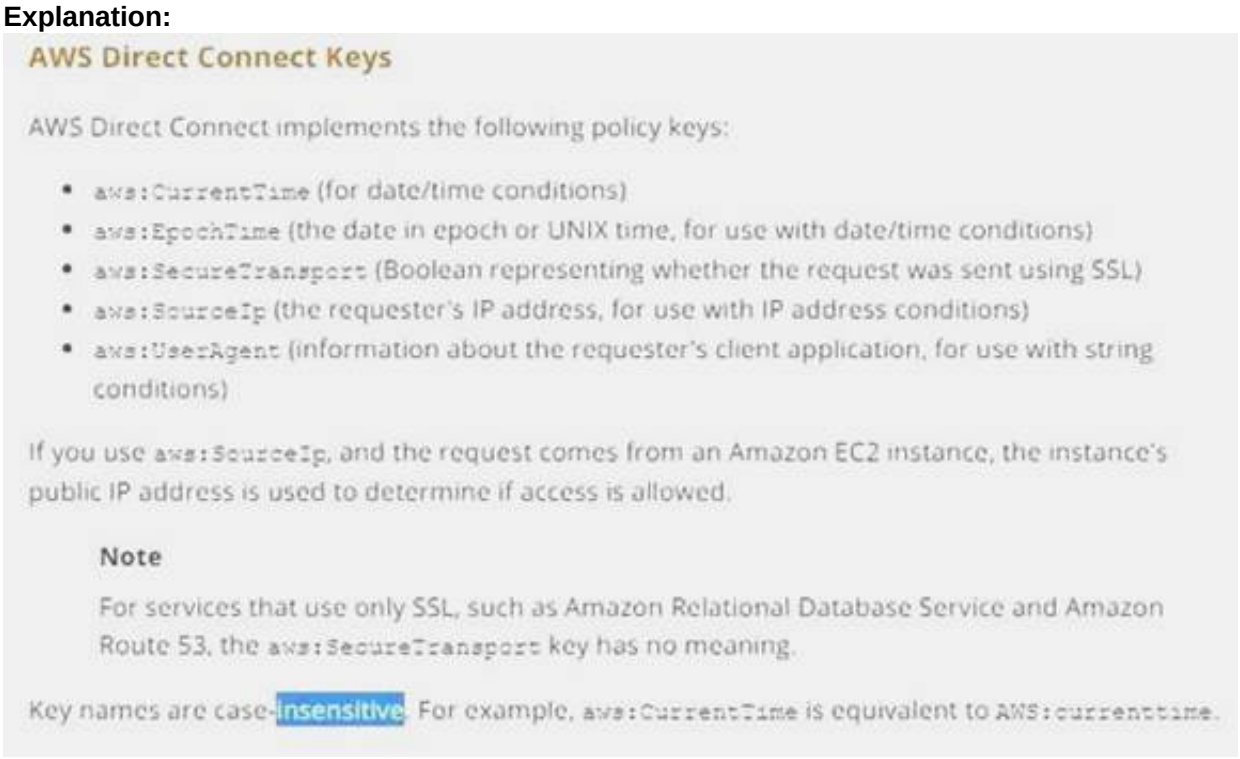

http://docs.aws.amazon.com/directconnect/latest/UserGuide/using_iam.html

NEW QUESTION 103
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

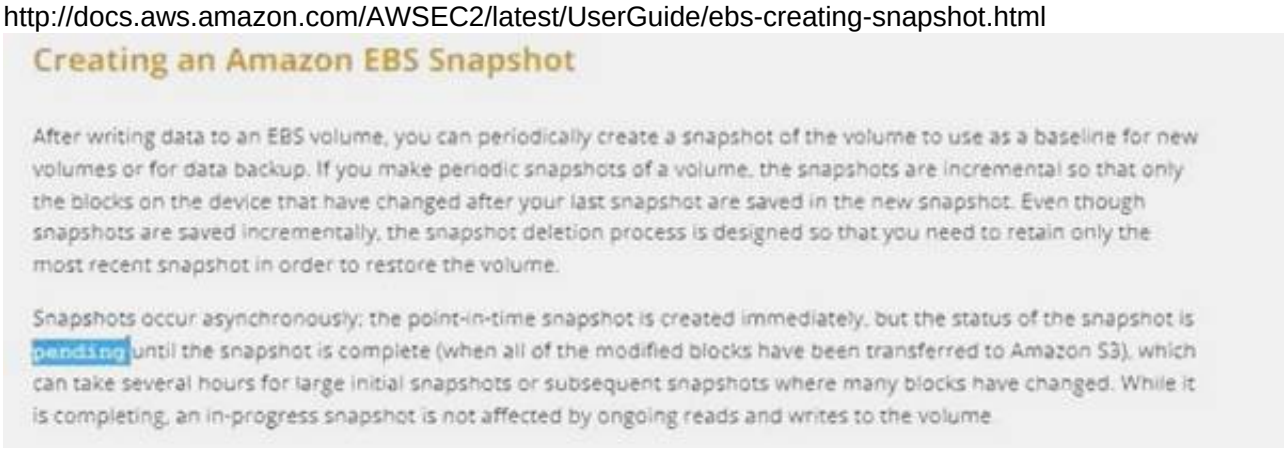

NEW QUESTION 106
What does the AWS Storage Gateway provide?

A. It allows to integrate on-premises IT environments with Cloud Storage.
B. A direct encrypted connection to Amazon S3.
C. It's a backup solution that provides an on-premises Cloud storage.
D. It provides an encrypted SSL endpoint for backups in the Clou

Answer: A

NEW QUESTION 108
Typically, you want to check your application whether a request generated an error before you spend any time processing results. The easiest way to find out if an error occurred is to look for an ____ node in the response from the Amazon RDS API.

A. Incorrect
B. Error
C. FALSE

Answer: B

Explanation:
Typically, you want your application to check whether a request generated an error before you spend any time processing results. The easiest way to find out if an error occurred is to look for an Error node in the response from the Amazon RDS API.
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/APITroubleshooting.html

NEW QUESTION 109
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 112
Amazon RDS DB snapshots and automated backups are stored in

A. Amazon S3
B. Amazon ECS Volume
C. Amazon RDS
D. Amazon EMR

Answer: A

NEW QUESTION 117
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 120
What are the Amazon EC2 API tools?

A. They don't exis
B. The Amazon EC2 AMI tools, instead, are used to manage permissions.
C. Command-line tools to the Amazon EC2 web service.
D. They are a set of graphical tools to manage EC2 instances.
E. They don't exis
F. The Amazon API tools are a client interface to Amazon Web Services.

Answer: B

NEW QUESTION 124
Disabling automated backups disable the point-in-time recovery.

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

Answer: C

NEW QUESTION 129
What is the maximum write throughput I can provision for a single Dynamic DB table?

A. 1,000 write capacity units
B. 100,000 write capacity units
C. Dynamic DB is designed to scale without limits, but if you go beyond 10,000 you have to contact AWS first.
D. 10,000 write capacity units

Answer: C

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

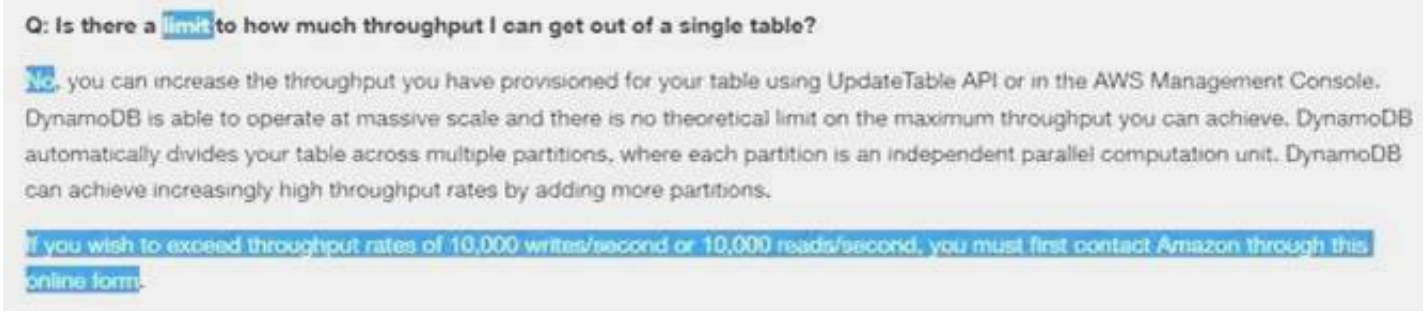

NEW QUESTION 131
Can I move a Reserved Instance from one Region to another?

A. No
B. Only if they are moving into GovCloud
C. Yes
D. Only if they are moving to US East from another region

Answer: A

NEW QUESTION 133
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 136
True or False: When you perform a restore operation to a point in time or from a DB Snapshot, a new DB Instance is created with a new endpoint.

A. FALSE
B. TRUE

Answer: B

Explanation:
Restoring From a DB Snapshot
Amazon RDS creates a storage volume snapshot of your DB instance, backing up the entire DB instance and not just individual databases. You can create a DB instance by restoring from this DB snapshot. When you restore the DB instance, you provide the name of the DB snapshot to restore from, and then provide a name for the new DB instance that is created from the restore. You cannot restore from a DB snapshot to an existing DB instance; a new DB instance is created when you restore. http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RestoreFromSnapshot.html

NEW QUESTION 137
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 141
While launching an RDS DB instance, on which page I can select the Availability Zone?

A. REVIEW
B. DB INSTANCE DETAILS
C. MANAGEMENT OPTIONS
D. ADDITIONAL CONFIGURATION

Answer: D

Explanation:
DB Instance detail -You just enable that your DB instance can be deploy in Multi-AZ. However,
you select the availability zone (Which AZ will be for primary and which one will be for secondary) in Additional configuration.

NEW QUESTION 142
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 146
While performing the volume status checks, if the status is insufficient-data, what does it mean?

A. the checks may still be in progress on the volume
B. the check has passed
C. the check has failed

Answer: A

Explanation:
If the status is insufficient-data, the checks may still be in progress on the volume. You can view the results of volume status checks to identify any impaired volumes and take any necessary actions.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring-volumestatus. html#monitoring-volume-checks

NEW QUESTION 147
By default, when an EBS volume is attached to a Windows instance, it may show up as any drive letter on the instance. You can change the settings of the ____ Service to set the drive letters of the EBS volumes per your specifications.

A. EBSConfig Service
B. AMIConfig Service
C. Ec2Config Service
D. Ec2-AMIConfig Service

Answer: C

Explanation:
Ec2Config Service is like sysprep and used specifically for windows instances. You can change parameters in OS before launching.

NEW QUESTION 148
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 149
When running my DB Instance as a Multi-AZ deployment, can I use the standby for read or write operations?

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

Answer: D

Explanation:

Q: When running my DB instance as a Multi-AZ deployment, can I use the standby for read or write operations?

No, the standby replica cannot serve read requests. Multi-AZ deployments are designed to provide enhanced database availability and durability, rather than read scaling benefits. As such, the feature uses synchronous replication between primary and standby. Our implementation makes sure the primary and the standby are constantly in sync, but precludes using the standby for read or write operations. If you are interested in a read scaling solution, please see the FAQs on Read Replicas.

NEW QUESTION 151
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 155
Amazon RDS automated backups and DB Snapshots are currently supported for only the _____ storage engine

A. MyISAM
B. InnoDB

Answer: B

NEW QUESTION 160
MySQL installations default to port _____ .

A. 3306
B. 443
C. 80
D. 1158

Answer: A

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

NEW QUESTION 161
Making your snapshot public shares all snapshot data with everyone. Can the snapshots with AWS Marketplace product codes be made public?

A. No
B. Yes

Answer: A

Explanation:
"Making your snapshot public shares all snapshot data with everyone; however, snapshots with AWS Marketplace product codes cannot be made public. Encrypted snapshots cannot be shared between
accounts or made public." http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-modifyingsnapshot- permissions.html
"This is not a valid option for encrypted snapshots or snapshots with AWS Marketplace product codes."

NEW QUESTION 165
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 166
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 168
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 172
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 173
What happens when you create a topic on Amazon SNS?

A. The topic is created, and it has the name you specified for it.
B. An ARN (Amazon Resource Name) is created.
C. You can create a topic on Amazon SQS, not on Amazon SNS.
D. This question doesn't make sens

Answer: B

NEW QUESTION 174
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 179
When automatic failover occurs, Amazon RDS will emit a DB Instance event to inform you that automatic failover occurred. You can use the ____ to return information about events related to your DB Instance

A. FetchFailure
B. DescriveFailure
C. DescribeEvents
D. FetchEvents

Answer: C

Explanation:
Q: Will I be alerted when automatic failover occurs?
Yes, Amazon RDS will emit a DB Instance event to inform you that automatic failover occurred. You can use the DescribeEvents to return information about events related to your DB Instance, or click the "DB Events" section of the AWS Management Console
https://aws.amazon.com/rds/faqs/

NEW QUESTION 182
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 185
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 187
Can we attach an EBS volume to more than one EC2 instance at the same time?

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

Answer: B

Explanation:
EBS is network attached storage that can only be attached to one instance at a time https://aws.amazon.com/ebs/getting-started/

NEW QUESTION 192
What does Amazon Route53 provide?

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

Answer: C

Explanation:
https://aws.amazon.com/route53/

NEW QUESTION 197
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 199
What is a Security Group?

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

Answer: D

Explanation:
A security group acts as a virtual firewall for your instance to control inbound and outbound traffic.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_SecurityGroups.html

NEW QUESTION 201
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 202
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 207
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 208
While controlling access to Amazon EC2 resources, which of the following acts as a firewall that controls the traffic allowed to reach one or more instances?

A. A security group
B. An instance type
C. A storage cluster
D. An object

Answer: A

Explanation:
A security group acts as a firewall that controls the traffic allowed to reach one or more instances. When you launch an instance, you assign it one or more security groups. http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/UsingIAM.html

NEW QUESTION 210
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 214
My Read Replica appears "stuck" after a Multi-AZ failover and is unable to obtain or apply updates from the source DB Instance. What do I do?

A. You will need to delete the Read Replica and create a new one to replace it.
B. You will need to disassociate the DB Engine and CK associate it.
C. The instance should be deployed to Single AZ and then moved to Multi- AZ once again
D. You will need to delete the DB Instance and create a new one to replace i

Answer: A

Explanation:
Q: My Amazon RDS for MySQL Read Replica appears "stuck" after a Multi-AZ failover and is unable to obtain or apply updates from the source DB Instance. What do I do? ... To resolve the current issue,
you will need to delete the Read Replica and create a new one to replace it. " https://aws.amazon.com/rds/faqs/

NEW QUESTION 216
In the 'Detailed' monitoring data available for your Amazon EBS volumes, Provisioned IOPS volumes automatically send _____ minute metrics to Amazon CloudWatch.

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

Answer: C

NEW QUESTION 217
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

NEW QUESTION 220
Can the string value of 'Key' be prefixed with laws?

A. No
B. Only for EC2 not S3
C. Yes
D. Only for S3 not EC

Answer: A

NEW QUESTION 225
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 229
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 234
Do the Amazon EBS volumes persist independently from the running life of an Amazon EC2 instance?

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

Answer: C

NEW QUESTION 237
What does Amazon RDS stand for?

A. Regional Data Server.
B. Relational Database Service.
C. Nothing.
D. Regional Database Servic

Answer: B

NEW QUESTION 240
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 243
What is the minimum time Interval for the data that Amazon CloudWatch receives and aggregates?

A. One second
B. Five seconds
C. One minute
D. Three minutes
E. Five minutes

Answer: C

Explanation:
Many metrics are received and aggregated at 1-minute intervals. Some are at 3-minute or 5-minute intervals.

NEW QUESTION 247
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 251
Location of Instances is _____

A. Regional
B. based on Availability Zone
C. Global

Answer: B

Explanation:
Regions and Availability Zones
Amazon EC2 is hosted in multiple locations world-wide. These locations are composed of regions and Availability Zones. Each region is a separate geographic area. Each region has multiple, isolated locations known as Availability Zones. Amazon EC2 provides you the ability to place resources, such as instances, and data in multiple locations. Resources aren't replicated across regions unless you do so specifically. http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availabilityzones. html#concepts-regions- availability-zones

```
$ aws ec2 describe-availability-zones --region us-east-1
{
    "AvailabilityZones": [
        {
            "State": "available",
            "RegionName": "us-east-1",
            "Messages": [],
            "ZoneName": "us-east-1b"
        },
        {
            "State": "available",
            "RegionName": "us-east-1",
            "Messages": [],
            "ZoneName": "us-east-1c"
        },
        {
            "State": "available",
            "RegionName": "us-east-1",
            "Messages": [],
            "ZoneName": "us-east-1d"
        }
    ]
}
```

NEW QUESTION 255
What does Amazon SES stand for?

A. Simple Elastic Server
B. Simple Email Service
C. Software Email Solution
D. Software Enabled Server

Answer: B

Explanation:
http://aws.amazon.com/ses/

> Amazon Simple Email Service (Amazon SES) is a cost-effective email service built on the reliable and scalable infrastructure that Amazon.com developed to serve its own customer base. With Amazon SES, you can send and receive email with no required minimum commitments – you pay as you go, and you only pay for what you use.

NEW QUESTION 258
Can I attach more than one policy to a particular entity?

A. Yes always
B. Only if within GovCloud
C. No
D. Only if within VPC

Answer: A

NEW QUESTION 261
Can I detach the primary (eth0) network interface when the instance is running or stopped?

A. Yes, You can.
B. N
C. You cannot
D. Depends on the state of the interface at the time

Answer: B

Explanation:
Each instance in a VPC has a default elastic network interface (the primary network interface, eth0) that is assigned a private IP address from the IP address range of your VPC. You cannot
detach a primary network interface from an instance.

NEW QUESTION 266
What's an ECU?

A. Extended Cluster User.
B. None of these.
C. Elastic Computer Usage.
D. Elastic Compute Uni

Answer: B

Explanation:
The EC2 Compute Unit (ECU) provides the relative measure of the integer processing power of an Amazon EC2 instance.
https://aws.amazon.com/ec2/faqs/

NEW QUESTION 269
What does the "Server Side Encryption" option on Amazon S3 provide?

A. It provides an encrypted virtual disk in the Cloud.
B. It doesn't exist for Amazon S3, but only for Amazon EC2.
C. It encrypts the files that you send to Amazon S3, on the server side.
D. It allows to upload files using an SSL endpoint, for a secure transfe

Answer: C

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

NEW QUESTION 274
Do the system resources on the Micro instance meet the recommended configuration for Oracle?

A. Yes completely
B. Yes but only for certain situations
C. Not in any circumstance

Answer: C

Explanation:
We recommend that you use db.t1.micro instances with Oracle to test setup and connectivity only; the system resources for a db.t1.micro instance do not meet the recommended configuration for Oracle. No Oracle options are supported on a db.t1.micro instance.
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuideConcepts.DBInstanceClass.html#Concep ts.DBInstanceCla ss.Previous

NEW QUESTION 279
To help you manage your Amazon EC2 instances, images, and other Amazon EC2 resources, you can assign your own metadata to each resource in the form of _____

A. special filters
B. functions
C. tags
D. wildcards

Answer: C

NEW QUESTION 280
Which choice is a storage option supported by Amazon EC2?

A. Amazon SNS store
B. Amazon Instance Store
C. Amazon AppStream store
D. None of these

Answer: B

Explanation:
Amazon EC2 supports the following storage options: Amazon Elastic Block Store (Amazon EBS) Amazon EC2 Instance Store
Amazon Simple Storage Service (Amazon S3) http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Storage.html

NEW QUESTION 283
Amazon EC2 provides a repository of public data sets that can be seamlessly integrated into AWS cloud-based applications. What is the monthly charge for using the public data sets?

A. A 1 time charge of 10$ for all the datasets.
B. 1$ per dataset per month
C. 10$ per month for all the datasets
D. There is no charge for using the public data sets

Answer: D

NEW QUESTION 286
Without _____ , you must either create multiple AWS accounts-each with its own billing and subscriptions to AWS products-or your employees must share the security credentials of a single AWS account.

A. Amazon RDS
B. Amazon Glacier
C. Amazon EMR
D. Amazon IAM

Answer: D

NEW QUESTION 287
Amazon RDS supports SOAP only through _____ .

A. HTTP or HTTPS
B. TCP/IP
C. HTTP
D. HTTPS

Answer: D

Explanation:
Amazon RDS supports SOAP only through HTTPS
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/using-soap-api.html


WSDL and Schema Definitions

You can access the Amazon Relational Database Service using the SOAP web services messaging protocol. This interface is described by a Web Services Description Language (WSDL) document, which defines the operations and security model for the particular service. The WSDL references an XML Schema document, which strictly defines the data types that might appear in SOAP requests and responses. For more information on WSDL and SOAP, see Web Services References.

Note

Amazon RDS supports SOAP only through HTTPS.

NEW QUESTION 291
Is creating a Read Replica of another Read Replica supported?

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

Answer: D

NEW QUESTION 296
What is the charge for the data transfer incurred in replicating data between your primary and standby?

A. Same as the standard data transfer charge
B. Double the standard data transfer charge
C. No charg
D. It is free
E. Half of the standard data transfer charge

Answer: C

Explanation:
Q: How much do Read Replicas cost? When does billing begin and end?
A Read Replica is billed as a standard DB Instance and at the same rates. Click here for more information on DB Instance billing visit this FAQ. Just like a standard DB Instance, the rate per “DB Instance hour” for a Read Replica is determined by the DB Instance class of the Read Replica –please see Amazon RDS detail page for up-to-date pricing. You are not charged for the data transfer incurred in replicating data between your source DB Instance and Read Replica.
Billing for a Read Replica begins as soon as the Read Replica has been successfully created (i.e. when status is listed as “active”). The Read Replica will continue being billed at standard Amazon RDS DB Instance hour rates until you issue a command to delete it.

NEW QUESTION 298
HTTP Query-based requests are HTTP requests that use the HTTP verb GET or POST and a Query parameter named____ .

A. Action
B. Value
C. Reset
D. Retrieve

Answer: A

Explanation:
http://docs.aws.amazon.com/AmazonS3/latest/dev/using-with-s3-actions.html

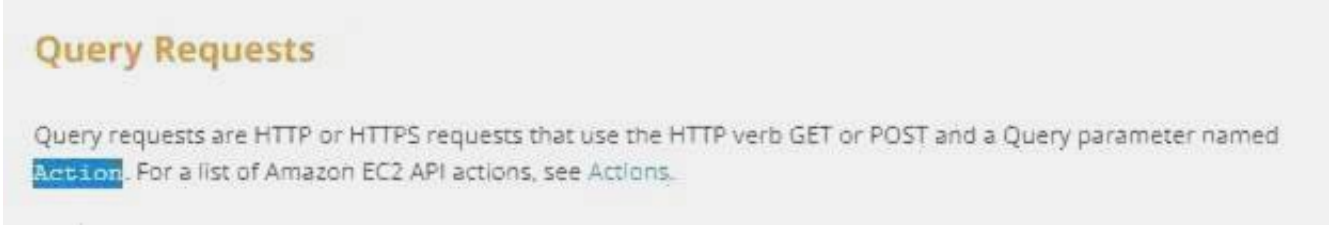
Query Requests

Query requests are HTTP or HTTPS requests that use the HTTP verb GET or POST and a Query parameter named Action. For a list of Amazon EC2 API actions, see Actions.

NEW QUESTION 300

_____ embodies the "share-nothing" architecture and essentially involves breaking a large database into several smaller databases. Common ways to split a database include 1) splitting tables that are not joined in the same query onto different hosts or 2) duplicating a table across multiple hosts and then using a hashing algorithm to determine which host receives a given update.

A. Sharding
B. Failure recovery
C. Federation
D. DDL operations

Answer: A

Explanation:
Sharding embodies the "share-nothing" architecture and essentially just involves breaking a larger database up into smaller databases. Common ways to split a database are:
Splitting tables that are not joined in the same query onto different hosts Duplicating a table across multiple hosts and then splitting where a row goes.
More detailed information on the pros and cons of sharing can be found at the following sites: http://technoroy.blogspot.com/2008/07/shard-database-design.html
http://www.hibernate.org/subprojects/shards.html
How Amazon RDS Helps With Sharing Maintenance Overhead

NEW QUESTION 304
When using consolidated billing there are two account types. What are they?

A. Paying account and Linked account
B. Parent account and Child account
C. Main account and Sub account.
D. Main account and Secondary accoun

Answer: A

Explanation:
You sign up for Consolidated Billing in the AWS Billing and Cost Management console, and designate your account as a payer account. Now your account can pay the charges of the other accounts, which are called linked accounts. The payer account and the accounts linked to it are called a Consolidated Billing account family. Source: http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/consolidated-billing.html

NEW QUESTION 309
You can modify the backup retention period; valid values are 0 (for no backup retention) to a maximum of ____ days.

A. 45
B. 35
C. 15
D. 5

Answer: B

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

NEW QUESTION 313
A Provisioned IOPS volume must be at least ____ GB in size

A. 1
B. 50
C. 20
D. 10

Answer: D

Explanation:
https://aws.amazon.com/ebs/details/

NEW QUESTION 316
After launching an instance that you intend to serve as a NAT (Network Address Translation) device in a public subnet you modify your route tables to have the NAT device be the target of internet bound traffic of your private subnet. When you try and make an outbound connection to the internet from an instance in the private subnet, you are not successful. Which of the following steps could resolve the issue?

A. Disabling the Source/Destination Check attribute on the NAT instance
B. Attaching an Elastic IP address to the instance in the private subnet
C. Attaching a second Elastic Network Interface (ENI) to the NAT instance, and placing it in the private subnet
D. Attaching a second Elastic Network Interface (ENI) to the instance in the private subnet, and placing it in the public subnet

Answer: A

Explanation:

NEW QUESTION 320
When using the following AWS services, which should be implemented in multiple Availability Zones for high availability solutions? Choose 2 answers

A. Amazon DynamoDB
B. Amazon Elastic Compute Cloud (EC2)
C. Amazon Elastic Load Balancing
D. Amazon Simple Notification Service (SNS)
E. Amazon Simple Storage Service (S3)

Answer: BC

NEW QUESTION 322
You have an EC2 Security Group with several running EC2 instances. You change the Security Group rules to allow inbound traffic on a new port and protocol, and launch several new instances in the same Security Group. The new rules apply:

A. Immediately to all instances in the security group.
B. Immediately to the new instances only.
C. Immediately to the new instances, but old instances must be stopped and restarted before the new rules apply.
D. To all instances, but it may take several minutes for old instances to see the changes.

Answer: A

Explanation:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html#vpc-securitygroups

NEW QUESTION 323
Which services allow the customer to retain full administrative privileges of the underlying EC2 instances? (Choose two.)

A. Amazon Relational Database Service
B. Amazon Elastic Map Reduce
C. Amazon ElastiCache
D. Amazon DynamoDB
E. AWS Elastic Beanstalk

Answer: BE

NEW QUESTION 328
A company is building a two-tier web application to serve dynamic transaction-based content. The data tier is leveraging an Online Transactional Processing (OLTP) database. What services should you leverage to enable an elastic and scalable web tier?

A. Elastic Load Balancing, Amazon EC2, and Auto Scaling
B. Elastic Load Balancing, Amazon RDS with Multi-AZ, and Amazon S3
C. Amazon RDS with Multi-AZ and Auto Scaling
D. Amazon EC2, Amazon DynamoDB, and Amazon S3

Answer: A

NEW QUESTION 329
You have decided to change the instance type for instances running in your application tier that is using Auto Scaling. In which area below would you change the instance type definition?

A. Auto Scaling policy
B. Auto Scaling group
C. Auto Scaling tags
D. Auto Scaling launch configuration

Answer: D

NEW QUESTION 333
When an EC2 EBS-backed (EBS root) instance is stopped, what happens to the data on any ephemeral store volumes?

A. Data is automatically saved in an EBS volume.
B. Data is unavailable until the instance is restarted.
C. Data will be deleted and will no longer be accessible.
D. Data is automatically saved as an EBS snapsho

Answer: C

Explanation:
When you stop a running instance, the following happens:
*The instance performs a normal shutdown and stops running; its status changes to stopping and then stopped.
*Any Amazon EBS volumes remain attached to the instance, and their data persists.
*Any data stored in the RAM of the host computer or the instance store volumes of the host computer is gone.

NEW QUESTION 335
You launch an Amazon EC2 instance without an assigned AWS identity and Access Management (IAM) role. Later, you decide that the instance should be running with an IAM role. Which action must you take in order to have a running Amazon EC2 instance with an IAM role assigned to it?

A. Create an image of the instance, and register the image with an IAM role assigned and an Amazon EBS volume mapping.
B. Create a new IAM role with the same permissions as an existing IAM role, and assign it to the running instance.
C. Create an image of the instance, add a new IAM role with the same permissions as the desired IAM role, and deregister the image with the new role assigned.
D. Create an image of the instance, and use this image to launch a new instance with the desired IAM role assigned.

Answer: D

NEW QUESTION 338
An instance is launched into a VPC subnet with the network ACL configured to allow all inbound traffic and deny all outbound traffic. The instance's security group is configured to allow SSH from any IP address and deny all outbound traffic. What changes need to be made to allow SSH access to the instance?

A. The outbound security group needs to be modified to allow outbound traffic.
B. The outbound network ACL needs to be modified to allow outbound traffic.
C. Nothing, it can be accessed from any IP address using SSH.
D. Both the outbound security group and outbound network ACL need to be modified to allow outbound traffic.

Answer: B

Explanation:
Need to open TCP Port 1024-65535 at Outbound Rules
"Allows outbound responses to the remote computer. Network ACLs are stateless, therefore this rule is required to allow response traffic for inbound requests."
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_ACLs.html

NEW QUESTION 341
For which of the following use cases are Simple Workflow Service (SWF) and Amazon EC2 an appropriate solution? (Choose two.)

A. Using as an endpoint to collect thousands of data points per hour from a distributed fileet of sensors
B. Managing a multi-step and multi-decision checkout process of an e-commerce website
C. Orchestrating the execution of distributed and auditable business processes
D. Using as an SNS (Simple Notification Service) endpoint to trigger execution of video transcoding jobs
E. Using as a distributed session store for your web application

Answer: BC

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

NEW QUESTION 343
What is a placement group?

A. A collection of Auto Scaling groups in the same region
B. A feature that enables EC2 instances to interact with each other via high bandwidth, low latency connections
C. A collection of authorized CloudFront edge locations for a distribution
D. A collection of Elastic Load Balancers in the same Region or Availability Zone

Answer: B

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

NEW QUESTION 345
You are configuring your company's application to use Auto Scaling and need to move user state information. Which of the following AWS services provides a

shared data store with durability and low latency?

A. AWS ElastiCache Memcached
B. Amazon Simple Storage Service
C. Amazon EC2 instance storage
D. Amazon DynamoDB

Answer: D

Explanation:
https://media.amazonwebservices.com/AWS_Storage_Options.pdf
To speed access to relevant data, many developers pair Amazon S3 with a database, such as Amazon DynamoDB or Amazon RDS. Amazon S3 stores the actual information, and the database serves as the repository for associated metadata (e.g., object name, size, keywords, and so on). Metadata in the database can easily be indexed and queried, making it very efficient to locate an object's reference
via a database query. This result can then be used to pinpoint and then retrieve the object itself from Amazon S3.

NEW QUESTION 348
Which of the following are characteristics of a reserved instance? (Choose three.)

A. It can be migrated across Availability Zones
B. It is specific to an Amazon Machine Image (AMI)
C. It can be applied to instances launched by Auto Scaling
D. It is specific to an instance Type
E. It can be used to lower Total Cost of Ownership (TCO) of a system

Answer: ACE

Explanation:
You can use Auto Scaling or other AWS services to launch the On-Demand instances that use your Reserved Instance benefits. For information about launching On-Demand instances, see Launch Your Instance. For information about launching instances using Auto Scaling, see the Auto Scaling User Guide.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts-on-demand-reservedinstances. html
https://forums.aws.amazon.com/thread.jspa?threadID=56501

NEW QUESTION 353
You run an ad-supported photo sharing website using S3 to serve photos to visitors of your site. At some point you find out that other sites have been linking to the photos on your site, causing loss to your business. What is an effective method to mitigate this?

A. Remove public read access and use signed URLs with expiry dates.
B. Use CloudFront distributions for static content.
C. Block the IPs of the offending websites in Security Groups.
D. Store photos on an EBS volume of the web serve

Answer: A

Explanation:
A signed URL includes additional information, for example, an expiration date and time, that gives you more control over access to your content.

NEW QUESTION 356
You are working with a customer who is using Chef configuration management in their data center. Which service is designed to let the customer leverage existing Chef recipes in AWS?

A. Amazon Simple Workflow Service
B. AWS Elastic Beanstalk
C. AWS CloudFormation
D. AWS OpsWorks

Answer: D

NEW QUESTION 358
When an EC2 instance that is backed by an S3-based AMI is terminated, what happens to the data on the root volume?

A. Data is automatically saved as an EBS snapshot.
B. Data is automatically saved as an EBS volume.
C. Data is unavailable until the instance is restarted.
D. Data is automatically delete

Answer: D

Explanation:
Using the legacy S3 based AMIs, either of the above terminates the instance and you lose all local and ephemeral storage (boot disk and /mnt) forever. Hope you remembered to save the important stuff elsewhere.

NEW QUESTION 360
You have a load balancer configured for VPC, and all back-end Amazon EC2 instances are in service. However, your web browser times out when connecting to the load balancer's DNS name. Which options are probable causes of this behavior? (Choose two.)

A. The load balancer was not configured to use a public subnet with an Internet gateway configured

B. The Amazon EC2 instances do not have a dynamically allocated private IP address
C. The security groups or network ACLs are not property configured for web traffic.
D. The load balancer is not configured in a private subnet with a NAT instance.
E. The VPC does not have a VGW configure

Answer: AC

Explanation:
There is no such thing as VGW. Hence E is not the correct answer.

NEW QUESTION 364
A company needs to deploy services to an AWS region which they have not previously used. The company currently has an AWS identity and Access Management (IAM) role for the Amazon EC2 instances, which permits the instance to have access to Amazon DynamoDB. The company wants their EC2 instances in the new region to have the same privileges. How should the company achieve this?

A. Create a new IAM role and associated policies within the new region
B. Assign the existing IAM role to the Amazon EC2 instances in the new region
C. Copy the IAM role and associated policies to the new region and attach it to the instances
D. Create an Amazon Machine Image (AMI) of the instance and copy it to the desired region using the AMI Copy feature

Answer: B

NEW QUESTION 367
You are deploying an application to collect votes for a very popular television show. Millions of users will submit votes using mobile devices. The votes must be collected into a durable, scalable, and highly available data store for real-time public tabulation. Which service should you use?

A. Amazon DynamoDB
B. Amazon Redshift
C. Amazon Kinesis
D. Amazon Simple Queue Service

Answer: A

Explanation:
This example looks at using AWS Lambda and Amazon API Gateway to build a dynamic voting application, which receives votes via SMS, aggregates the totals into Amazon DynamoDB, and uses Amazon Simple Storage Service (Amazon S3) to display the results in real time.
http://www.allthingsdistributed.com/2016/06/aws-lambda-serverless-reference-architectures.html

NEW QUESTION 371
The Trusted Advisor service provides insight regarding which four categories of an AWS account?

A. Security, fault tolerance, high availability, and connectivity
B. Security, access control, high availability, and performance
C. Performance, cost optimization, security, and fault tolerance
D. Performance, cost optimization, access control, and connectivity

Answer: C

Explanation:

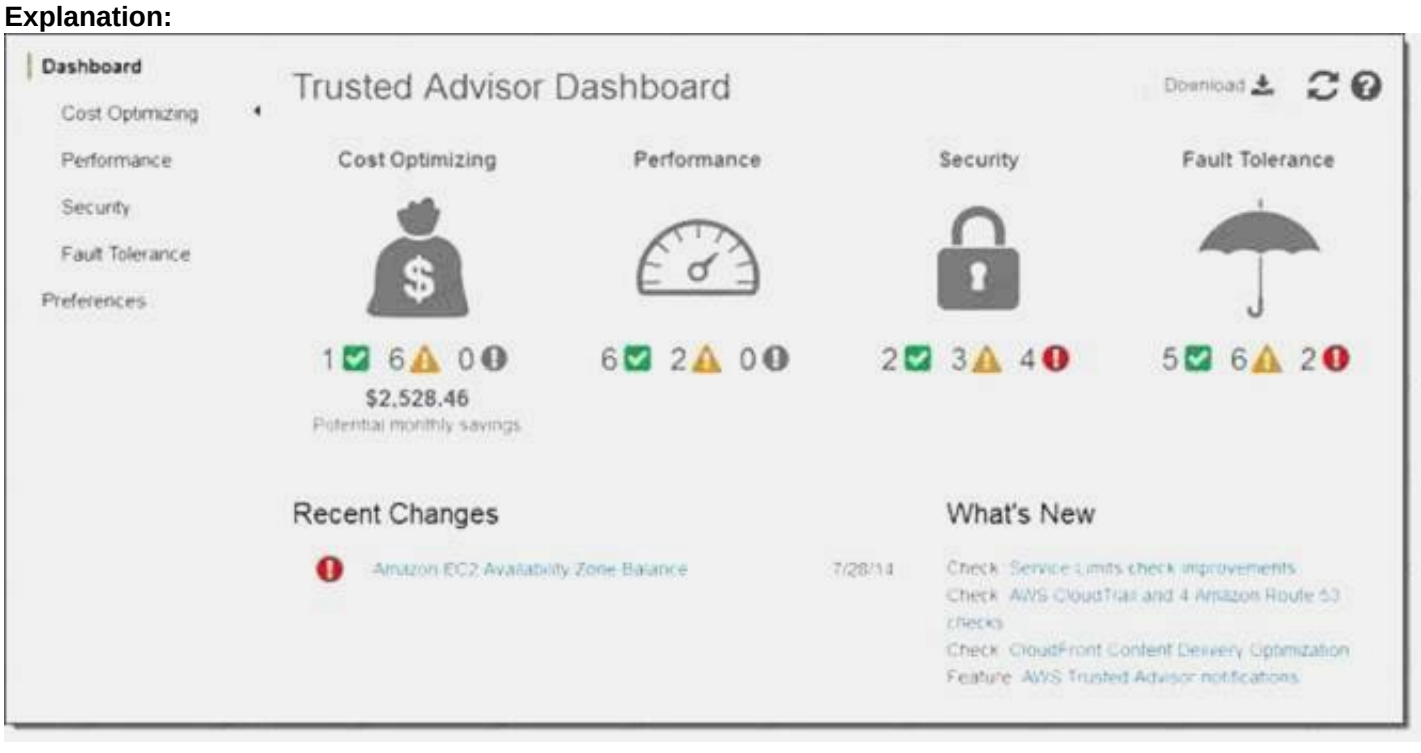

NEW QUESTION 376
A photo-sharing service stores pictures in Amazon Simple Storage Service (S3) and allows application sign-in using an OpenID Connect-compatible identity provider. Which AWS Security Token Service approach to temporary access should you use for the Amazon S3 operations?

A. SAML-based Identity Federation
B. Cross-Account Access
C. AWS Identity and Access Management roles

D. Web Identity Federation

Answer: D

Explanation:
Web identity federation - You can let users sign in using a well-known third party identity provider such as Login with Amazon, Facebook, Google, or any OpenID Connect (OIDC) 2.0 compatible provider. AWS STS web identity federation supports Login with Amazon, Facebook, Google, and any OpenID Connect (OICD)-compatible identity provider.

NEW QUESTION 381
A customer wants to track access to their Amazon Simple Storage Service (S3) buckets and also use this information for their internal security and access audits. Which of the following will meet the Customer requirement?

A. Enable AWS CloudTrail to audit all Amazon S3 bucket access.
B. Enable server access logging for all required Amazon S3 buckets.
C. Enable the Requester Pays option to track access via AWS Billing
D. Enable Amazon S3 event notifications for Put and Pos

Answer: B

Explanation:
References:
http://docs.aws.amazon.com/AmazonS3/latest/dev/ServerLogs.html http://docs.aws.amazon.com/AmazonS3/latest/dev/cloudtrail-logging.html

NEW QUESTION 386
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual SAA-C01 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the SAA-C01 Product From:

https://www.2passeasy.com/dumps/SAA-C01/

Money Back Guarantee

SAA-C01 Practice Exam Features:

* SAA-C01 Questions and Answers Updated Frequently

* SAA-C01 Practice Questions Verified by Expert Senior Certified Staff

* SAA-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* SAA-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year