Vendor: Amazon

Exam Code: AWS Certified Solutions Architect - Professional

Professional

Version: 17.121



PASS YOUR EXAM
IN FIRST ATTEMPT
REAL EXAM QUESTIONS

QUESTION 1
You've been brought in as solutions architect to assist an enterprise customer with their migration of an e-commerce platform to Amazon Virtual Private Cloud (VPC).
The previous architect has already deployed a 3-tier VPC.
The configuration is as follows:

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

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

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

Answer: A
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.

QUESTION 2
You are designing the network infrastructure for an application server in Amazon VPC. Users will access all the application instances from the Internet, as well as from an on-premises network. The on-premises network is connected to your VPC over an AWS Direct Connect link.
How would you design routing to meet the above requirements?

A.	Configure a single routing table with a default route via the Internet gateway. Propagate a default route via BGP on the AWS Direct Connect customer router. Associate the routing table with all VPC subnets.
B.	Configure a single routing table with a default route via the Internet gateway. Propagate specific routes for the on-premises networks via BGP on the AWS Direct Connect

	customer router. Associate the routing table with all VPC subnets.
C.	Configure two routing tables: one that has a default route via the Internet gateway, and another that has a default route via the VPN gateway. Associate both routing tables with each VPC subnet.
D.	Configure a single routing table with two default routes: one to the Internet via an Internet gateway, the other to the on-premises network via the VPN gateway. Use this routing table across all subnets in your VPC.

Answer: B

QUESTION 3
Within the IAM service a GROUP is regarded as a:

A.	A collection of AWS accounts
B.	It's the group of EC2 machines that gain the permissions specified in the GROUP.
C.	There's no GROUP in IAM, but only USERS and RESOURCES.
D.	A collection of users.

Answer: D
Explanation:
Use groups to assign permissions to IAM users
Instead of defining permissions for individual IAM users, it's usually more convenient to create groups that relate to job functions (administrators, developers, accounting, etc.), define the relevant permissions for each group, and then assign IAM users to those groups. All the users in an IAM group inherit the permissions assigned to the group. That way, you can make changes for everyone in a group in just one place. As people move around in your company, you can simply change what IAM group their IAM user belongs to.
http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html#use-groups-for-permissions

QUESTION 4
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

QUESTION 5
In the Amazon RDS Oracle DB engine, the Database Diagnostic Pack and the Database Tuning Pack are only available with _____

A.	Oracle Standard Edition
B.	Oracle Express Edition
C.	Oracle Enterprise Edition

D.	None of these

Answer: C
Explanation:
https://www.pythian.com/blog/a-most-simple-cloud-is-amazon-rds-for-oracle-right-for-you/

QUESTION 6
You have deployed a web application, targeting a global audience across multiple AWS Regions under the domain name example.com.
You decide to use Route53 Latency-Based Routing to serve web requests to users from the region closest to the user. To provide business continuity in the event of server downtime you configure weighted record sets associated with two web servers in separate Availability Zones per region. During a DR test you notice that when you disable all web servers in one of the regions Route53 does not automatically direct all users to the other region.
What could be happening? Choose 2 answers

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

Answer: AE
Explanation:
http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-complex-configs.html
For both latency alias resource record sets, you set the value of “Evaluate Target Health” to Yes.
You use the Evaluate Target Health setting for each latency alias resource record set to make Amazon Route 53 evaluate the health of the alias targets—the weighted resource record sets—and respond accordingly.

QUESTION 7
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 the 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 day, create a "LastUpdated" attribute in your DynamoDB table that would represent the timestamp of the last update and use it as a filter
B.	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
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.	Send also each write into an SQS queue in the second region, use an auto-scaling group behind

	the SQS queue to replay the write in the second region

Answer: B
Explanation:
Export of dynamo DB is incremental and it will amend the backup with latest changes.

QUESTION 8
Your company currently has a 2-tier web application running in an on-premises data center. You have experienced several infrastructure failures in the past few months resulting in significant financial losses. Your CIO is strongly considering moving the application to AWS. While working on achieving buy-In from the other company executives, he asks you to develop a disaster recovery plan to help improve business continuity in the short term. He specifies a target Recovery Time Objective (RTO) of 4 hours and a Recovery Point Objective (RPO) of 1 hour or less. He also asks you to implement the solution within 2 weeks.
Your database is 200GB in size and you have a 20Mbps Internet connection.
How would you do this while minimizing costs?

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

Answer: A

QUESTION 9
You would like to create a mirror image of your production environment in another region for disaster recovery purposes.
Which of the following AWS resources do not need to be recreated in the second region? Choose 2 answers

A.	Route S3 Record Sets
B.	Launch Configurations
C.	EC2 Key Pairs
D.	Security Groups
E.	IAM Roles
F.	Elastic IP Addresses (EIP)

Answer: AE
Explanation:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/resources.html
As per the document defined, new IPs should be reserved not the same ones.

Elastic IP Addresses are static IP addresses designed for dynamic cloud computing. Unlike traditional static IP addresses, however, Elastic IP addresses enable you to mask instance or Availability Zone failures by programmatically remapping your public IP addresses to instances in your account in a particular region. For DR, you can also pre-allocate some IP addresses for the most critical systems so that their IP addresses are already known before disaster strikes. This can simplify the execution of the DR plan.

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, then 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 base 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 RDS 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 SES to send emails to customers.
C.	Use an SQS queue to manage all process tasks. Use an Auto Scaling group of EC2 instances that poll the tasks and execute them. Use SES to send emails to customers.
D.	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.

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

QUESTION 11
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 A2 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 a Multi-AZ RDS (Relational Database Service) deployment.
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 in each AZ inside an Auto Scaling Group behind an ELB, 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 ELB, and one RDS (Relational Database Service) instance deployed with read replicas in the two other AZs.

Answer: B
Explanation:
http://highscalability.com/blog/2016/1/11/a-beginners-guide-to-scaling-to-11-million-users-on-amazons.html
https://www.airpair.com/aws/posts/building-a-scalable-web-app-on-amazon-web-services-p1

QUESTION 12
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 of the application with the anticipated additional load? Why?

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

Answer: A
Explanation:
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

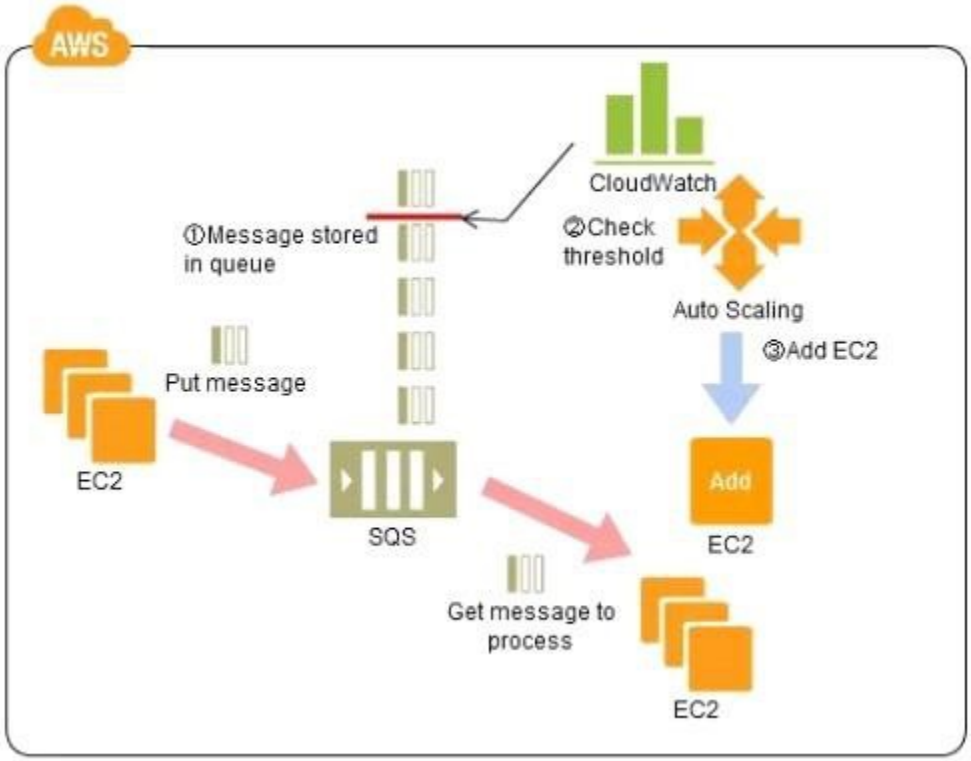
To mitigate the impact of an availability zone failure, locate your nodes in as many availability zones as possible. In the unlikely event of an AZ failure, you will lose only the data cached in that AZ, not the data cached in the other AZs.

QUESTION 13
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 15 minute DB backups stored in Glacier 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 hourly DB backups to S3, with transaction logs stored in S3 every 5 minutes.

Answer: D

QUESTION 14
Refer to the architecture diagram above of a batch processing solution using Simple Queue Service (SQS) to set up a message queue between EC2 instances which are used as batch processors. CloudWatch monitors the number of job requests (queued messages) and an Auto Scaling group adds or deletes batch servers automatically based on parameters set in CloudWatch alarms.
You can use this architecture to implement which of the following features in a cost effective and efficient manner?



A.	Coordinate number of EC2 instances with number of Job requests automatically, thus improving

	cost effectiveness.
B.	Reduce the overall time for executing Jobs through parallel processing by allowing a busy EC2 instance that receives a message to pass it to the next instance in a daisy-chain setup.
C.	Implement fault tolerance against EC2 instance failure since messages would remain in SQS and work can continue with recovery of EC2 instances. Implement fault tolerance against SQS failure by backing up messages to S3.
D.	Handle high priority Jobs before lower priority Jobs by assigning a priority metadata field to SQS messages.
E.	Implement message passing between EC2 instances within a batch by exchanging messages through SQS.

Answer: A
Explanation:
https://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-using-sqs-queue.html

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

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

QUESTION 16
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 filesystems 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 database to S3 using Oracle RMAN. Backup the EC2 instances using AMIs, and supplement with EBS snapshots for individual volume

	restore.
C.	Backup RDS using a Multi-AZ Deployment. Backup the EC2 instances using AMIs, and supplement by copying filesystem data to S3 to provide file level restore.
D.	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.

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.

QUESTION 17
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 four EBS volumes are configured as a single RAID 0 device, and each Provisioned IOPS volume is provisioned with 4,000 IOPS (4,000 16KB reads or writes), for a total of 16,000 random IOPS on the instance. The EC2 instance initially delivers the expected 16,000 IOPS random read and write performance. Sometime later, in order to increase the total random I/O performance of the instance, you add an additional two 500 GB EBS Provisioned IOPS volumes to the RAID. Each volume is provisioned to 4,000 IOPs like the original four, for a total of 24,000 IOPS on the EC2 instance. Monitoring shows that the EC2 instance CPU utilization increased from 50% to 70%, but the total random IOPS measured at the instance level does not increase at all.
What is the problem and a valid solution?

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

Answer: C

QUESTION 18
A customer has a 10 GB AWS Direct Connect connection to an AWS region where they have a web application hosted on Amazon Elastic Computer Cloud (EC2). The application has dependencies on an on-premises mainframe database that uses a BASE (Baste Available, Soft state, Eventual consistency) rather than an ACID (Atomicity, Consistency, Isolation, Durability) consistency model. The application is exhibiting undesirable behavior because the database is not able to handle the volume of writes. How can you reduce the load on your on-premises database resources in the most cost-effective way?

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

Answer: C
Explanation:
https://aws.amazon.com/sqs/faqs/

QUESTION 19
Your company plans to host a large donation website on Amazon Web Services (AWS).
You anticipate a large and undetermined amount of traffic that will create many database writes.
To be certain that you do not drop any writes to a database hosted on AWS, which service should you use?

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

Answer: A
Explanation:
https://aws.amazon.com/sqs/faqs/
There is no limit on the number of messages that can be pushed onto SQS. The retention period of the SQS is 4 days by default and it can be changed to 14 days. This will make sure that no writes are missed.

QUESTION 20
You have been asked to design the storage layer for an application. The application requires disk performance of at least 100,000 IOPS. In addition, the storage layer must be able to survive the loss of an individual disk, EC2 instance, or Availability Zone without any data loss.
The volume you provide must have a capacity of at least 3 TB.
Which of the following designs will meet these objectives?

A.	Instantiate a c3 8xlarge instance in us-east-1. Provision 4x1TB EBS volumes, attach them to the instance, and configure them as a single RAID 5 volume. Ensure that EBS snapshots are performed every 15 minutes.
B.	Instantiate a c3 8xlarge instance in us-east-1. Provision 3xlTB EBS volumes, attach them to the Instance, and configure them as a single RAID 0 volume. Ensure that EBS snapshots are performed every 15 minutes.
C.	Instantiate an 12 8xlarge instance in us-east-1a. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instance. Provision 3x1TB EBS volumes, attach them to the instance, and configure them as a second RAID 0 volume.

	Configure synchronous, block-level replication from the ephemeral-backed volume to the EBS-backed volume.
D.	Instantiate a c3 8xlarge instance in us-east-1. Provision an AWS Storage Gateway and configure it for 3 TB of storage and 100,000 IOPS. Attach the volume to the instance.
E.	Instantiate an 12 8xlarge instance in us-east-1a. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instance. Configure synchronous, block-level replication to an identically configured instance in us-east-1b.

Answer: E
Explanation:
It doesn't protect against a loss of two EC2 instances or two AZs, but the question asks about protection of ONE disk, EC2 instance or AZ loss.
https://acloud.guru/course/aws-certified-solutions-architect-associate/discuss/-KJdi4tFMp2x_O88J6U4/an-architecture-design-question

QUESTION 21
Your company has HQ in Tokyo and branch offices all over the world and is using a logistics software with a multi-regional deployment on AWS in Japan, Europe and US.
The logistic software has a 3-tier architecture and currently uses MySQL 5.6 for data persistence.
Each region has deployed its own database.
In the HQ region you run an hourly batch process reading data from every region to compute cross- regional reports that are sent by email to all offices.
This batch process must be completed as fast as possible to quickly optimize logistics.
How do you build the database architecture in order to meet the requirements?

A.	For each regional deployment, use MySQL on EC2 with a master in the region and use S3 to copy data files hourly to the HQ region.
B.	For each regional deployment, use RDS MySQL with a master in the region and send hourly RDS snapshots to the HQ region.
C.	Use Direct Connect to connect all regional MySQL deployments to the HQ region and reduce network latency for the batch process.
D.	For each regional deployment, use RDS MySQL with a master in the region and a read replica In the HQ region.
E.	For each regional deployment, use MySQL on EC2 with a master in the region and send hourly EBS snapshots to the HQ region.

Answer: D

QUESTION 22
Company B is launching a new game app for mobile devices. Users will log into the game using their existing social media account. To streamline data capture, Company B would like to directly save player data and scoring information from the mobile app to a DynamoDB table named ScoreData. When a user saves their game, the progress data will be stored to the GameState S3 bucket. What is the best approach for storing data to DynamoDB and S3?

A.	Use Login with Amazon allowing users to sign in with an Amazon account providing the mobile app with access to the ScoreData DynamoDB table and the GameState S3 bucket.
B.	Use temporary security credentials that assume a role providing access to the ScoreData DynamoDB table and the GameState S3 bucket using web identity federation
C.	Use an IAM user with access credentials assigned a role providing access to the ScoreData DynamoDB table and the GameState S3 bucket for distribution with the mobile app

D.	Use an EC2 instance that is launched with an EC2 role providing access to the ScoreData DynamoDB table and the GameState S3 bucket that communicates with the mobile app via web services

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.”

QUESTION 23
You have recently joined a startup company building sensors to measure street noise and air quality in urban areas.
The company has been running a pilot deployment of around 100 sensors for 3 months. Each sensor uploads 1KB of sensor data every minute to a backend hosted on AWS. During the pilot, you measured a peak of 10 IOPS on the database, and you stored an average of 3GB of sensor data per month in the database.
The current deployment consists of a load-balanced, auto scaled Ingestion layer using EC2 instances, and a PostgreSQL RDS database with 500GB standard storage
The pilot is considered a success and your CEO has managed to get the attention of some potential Investors. The business plan requires a deployment of at least 100k sensors which needs to be supported by the backend.
You also need to store sensor data for at least two years to be able to compare year over year improvements.
To secure funding, you have to make sure that the platform meets these requirements and leaves room for further scaling.
Which setup will meet the requirements?

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

Answer: C
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 autoscaling group.
Also, this example solution from AWS is some what similar for reference.
http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_timeseriesprocessing_16.pdf

QUESTION 24
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 client. 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.
B.	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 for each customer with a Bucket Policy that permits access only to that one customer.
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 threshold. Load a central list of FTP users from S3 as part of the User Data startup script on each instance.

Answer: A
Explanation:
https://aws.amazon.com/blogs/security/writing-iam-policies-grant-access-to-user-specific-folders-in-an-amazon-s3-bucket/

QUESTION 25
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 will meet the initial requirements for the collection platform?

A.	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.
B.	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.
C.	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.
D.	Utilize EMR to collect the inbound sensor data, analyze the data from EMR with Amazon Kinesis and save the results to DynamoDB.

Answer: A

QUESTION 26
You need a persistent and durable storage to trace call activity of an IVR (Interactive Voice Response) system. Call duration is mostly in the 2-3 minutes timeframe. Each traced call can be either active or terminated. An external application needs to know each minute the list of currently active calls. Usually there are a few calls/second, but once per month there is a periodic peak up to 1000 calls/second for a few hours. The system is open 24/7 and any downtime should be

avoided. Historical data is periodically archived to files. Cost saving is a priority for this project.
What database implementation would better fit this scenario, keeping costs as low as possible?

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

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

QUESTION 27
A 3-Ber e-commerce web application is currently deployed on-premises, and will be migrated to AWS for greater scalability and elasticity. The web tier currently shares read-only data using a network distributed file system. The app server tier uses a clustering mechanism for discovery and shared session state that depends on IP multicast. The database tier uses shared-storage clustering to provide database failover capability, and uses several read slaves for scaling. Data on all servers and the distributed file system directory is backed up weekly to off-site tapes.
Which AWS storage and database architecture meets the requirements of the application?

A.	Web servers: store read-only data in S3, and copy from S3 to root volume at boot time. App servers: share state using a combination of DynamoDB and IP unicast. Database: use RDS with multi-AZ deployment and one or more read replicas. Backup: web servers, app servers, and database backed up weekly to Glacier using snapshots.
B.	Web servers: store read-only data in an EC2 NFS server, mount to each web server at boot time. App servers: share state using a combination of DynamoDB and IP multicast. Database: use RDS with multi- AZ deployment and one or more Read Replicas. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.
C.	Web servers: store read-only data in S3, and copy from S3 to root volume at boot time. App servers: share state using a combination of DynamoDB and IP unicast. Database: use RDS with multi-AZ deployment and one or more Read Replicas. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.
D.	Web servers: store read-only data in S3, and copy from S3 to root volume at boot time.

App servers: share state using a combination of DynamoDB and IP unicast. Database: use RDS with multi-AZ deployment. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.

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

QUESTION 28
You require the ability to analyze a large amount of data which is stored on Amazon S3 using Amazon Elastic MapReduce.
You are using the cc2.8xlarge instance type, whose CPUs are mostly idle during processing.
Which of the below would be the most cost efficient way to reduce the runtime of the job?

A.	Create fewer, larger files m Amazon S3.
B.	Use smaller instances that have higher aggregate I/O performance.
C.	Create more, smaller files on Amazon S3.
D.	Add additional cc2.8xlarge instances by introducing a task group.

Answer: B

QUESTION 29
You are the new IT architect in a company that operates a mobile sleep tracking application.
When activated at night, the mobile app is sending collected data points of 1 kilobyte every 5 minutes to your backend.
The backend takes care of authenticating the user and writing the data points into an Amazon DynamoDB table.
Every morning, you scan the table to extract and aggregate last night's data on a per user basis, and store the results in Amazon S3. Users are notified via Amazon SNS mobile push notifications that new data is available, which is parsed and visualized by the mobile app. Currently you have around 100k users who are mostly based out of North America.
You have been tasked to optimize the architecture of the backend system to lower cost.
What would you recommend? Choose 2 answers

A.	Have the mobile app access Amazon DynamoDB directly Instead of JSON files stored on Amazon S3.
B.	Write data directly into an Amazon Redshift cluster replacing both Amazon DynamoDB and Amazon S3.
C.	Introduce an Amazon SQS queue to buffer writes to the Amazon DynamoDB table and reduce provisioned write throughput.
D.	Introduce Amazon Elasticache to cache reads from the Amazon DynamoDB table and reduce provisioned read throughput.
E.	Create a new Amazon DynamoDB table each day and drop the one for the previous day after its

	data is on Amazon S3.

Answer: CD
Explanation:
https://d0.awsstatic.com/whitepapers/performance-at-scale-with-amazon-elasticache.pdf

QUESTION 30
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 jobs. Use Reserved Instances for Amazon Redshift.
B.	Use reduced redundancy storage (RRS) for PDF and .csv data in S3. Add Spot Instances to EMR jobs. Use Spot Instances for Amazon Redshift.
C.	Use reduced redundancy storage (RRS) for PDF and .csv data In Amazon S3. Add Spot Instances to Amazon EMR jobs. Use Reserved Instances for Amazon Redshift.
D.	Use reduced redundancy storage (RRS) for all data in Amazon S3. Add Spot Instances to Amazon EMR jobs. Use Reserved Instances for Amazon Redshift.

Answer: C
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/

QUESTION 31
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 If 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 queue. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few days. CloudFront to serve HLS transcoded videos from EC2.
B.	Elastic Transcoder to transcode original high-resolution MP4 videos to HLS. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few days. CloudFront to serve HLS transcoded videos from EC2.
C.	Elastic Transcoder to transcode original high-resolution MP4 videos to HLS. S3 to host videos with Lifecycle Management to archive original files to Glacier after a few days. CloudFront to serve HLS transcoded videos from S3.
D.	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 queue. S3 to host videos with Lifecycle Management to archive all files to Glacier after a few days. CloudFront to serve HLS transcoded videos from Glacier.

Answer: C

QUESTION 32
You deployed your company website using Elastic Beanstalk and you enabled log file rotation to S3. An Elastic MapReduce 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 CloudFront 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.	Change your log collection process to use CloudWatch ELB metrics as input of the Elastic MapReduce Job.
B.	Turn on CloudTrail and use trail log files on S3 as input of the Elastic MapReduce job.
C.	Enable CloudFront to deliver access logs to S3 and use them as input of the Elastic MapReduce job.
D.	Use Elastic Beanstalk "Restart App Server(s)" option to update log delivery to the Elastic MapReduce job.
E.	Use Elastic Beanstalk "Rebuild Environment" option to update log delivery to the Elastic MapReduce job.

Answer: C
Explanation:
http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/AccessLogs.html

QUESTION 33
A web company is looking to implement an intrusion detection and prevention system into their deployed VPC. This platform should have the ability to scale to thousands of instances running inside of the VPC.
How should they architect their solution to achieve these goals?

A.	Configure each host with an agent that collects all network traffic and sends that traffic to the IDS/IPS platform for inspection.
B.	Configure an instance with monitoring software and the elastic network interface (ENI) set to promiscuous mode packet sniffing to see all traffic across the VPC.

C.	Create a second VPC and route all traffic from the primary application VPC through the second VPC where the scalable virtualized IDS/IPS platform resides.
D.	Configure servers running in the VPC using the host-based "route" commands to send all traffic through the platform to a scalable virtualized IDS/IPS.

Answer: C

QUESTION 34
You are running a successful multitier web application on AWS and your marketing department has asked you to add a reporting tier to the application. The reporting tier will aggregate and publish status reports every 30 minutes from user-generated information that is being stored in your web application's database.
You are currently running a Multi-AZ RDS MySQL instance for the database tier.
You also have implemented ElastiCache as a database caching layer between the application tier and database tier. Please select the answer that will allow you to successfully implement the reporting tier with as little impact as possible to your database:

A.	Launch a RDS Read Replica connected to your Multi AZ master database and generate reports by querying the Read Replica.
B.	Continually send transaction logs from your master database to an S3 bucket and generate the reports off the S3 bucket using S3 byte range requests.
C.	Generate the reports by querying the ElastiCache database caching tier.
D.	Generate the reports by querying the synchronously replicated standby RDS MySQL instance maintained through Multi-AZ.

Answer: A

QUESTION 35
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 batch process this data and used RabbitMQ, 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 SNS to pass job messages, use CloudWatch alarms to terminate spot worker instances when they become idle. Once data is processed, change the storage class of the S3 object to Glacier.
B.	Use SQS for passing job messages, use CloudWatch alarms to terminate EC2 worker instances when they become idle. 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 SQS. Once data is processed, change the storage class of the S3 objects to Reduced Redundancy Storage.
D.	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.

Answer: D
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.

QUESTION 36

You're running an application on-premises due to its dependency on non-x86 hardware and want to use AWS for data backup. Your backup application is only able to write to POSIX-compatible, block-based storage. You have 140TB of data and would like to mount it as a single folder on your file server. Users must be able to access portions of this data while the backups are taking place. What backup solution would be most appropriate for this use case?

A.	Use Storage Gateway and configure it to use Gateway Cached volumes
B.	Use Storage Gateway and configure it to use Gateway Stored volumes
C.	Configure your backup software to use S3 as the target for your data backups
D.	Configure your backup software to use Glacier as the target for your data backups

Answer: D

Explanation:

Rejecting other 3 options as data stored to S3 and S3 is object storage, uses a flat namespace and isn’t meant to serve as a standalone, POSIX-compliant file system. VTS shelf would have been a better option, but the question might be old. Also unclear for the need to take care of users been able to access portions of data.

QUESTION 37

A corporate web application is deployed within an Amazon Virtual Private Cloud (VPC), and is connected to the corporate data center via an IPsec VPN. The application must authenticate against the on- premises LDAP server. After authentication, each logged-in user can only access an Amazon Simple Storage Space (S3) keyspace specific to that user.

Which two approaches can satisfy these objectives? Choose 2 answers

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

	the appropriate S3 bucket.

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

QUESTION 38
You are designing a multi-platform web application for AWS. The application will run on EC2 instances and will be accessed from PCs, tablets and smart phones, supported accessing platforms are Windows, MacOS, IOS and Android. Separate sticky session and SSL certificate setups are required for different platform types. Which of the following describes the most cost effective and performance efficient architecture setup?

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

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

QUESTION 39
You are tasked with moving a legacy application from a virtual machine running inside your datacenter to an Amazon VPC. Unfortunately, this app requires access to a number of on-premises services and no one who configured the app still works for your company. Even worse, there's no documentation for it. What will allow the application running inside the VPC to reach back and access its internal dependencies without being reconfigured? Choose 3 answers

A.	A VM Import of the current virtual machine

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

Answer: ADF

QUESTION 40
You are designing a connectivity solution between on-premises infrastructure and Amazon VPC.
Your servers on-premises will be communicating with your VPC instances.
You will be establishing IPsec tunnels over the Internet.
You will be using VPN gateways, and terminating the IPsec tunnels on AWS supported customer gateways.
Which of the following objectives would you achieve by implementing an IPsec tunnel as outlined above? Choose 4 answers

A.	Peer identity authentication between VPN gateway and customer gateway.
B.	End-to-end identity authentication.
C.	Data integrity protection across the Internet.
D.	End-to-end protection of data in transit.
E.	Data encryption across the Internet.
F.	Protection of data in transit over the Internet.

Answer: ACEF

QUESTION 41
You are designing an intrusion detection/prevention (IDS/IPS) solution for a customer web application in a single VPC.
You are considering the options for Implementing IDS/IPS protection for traffic coming from the Internet.
Which of the following options would you consider? Choose 2 answers

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

Answer: AC
Explanation:
EC2 does not allow promiscuous mode, and you cannot put something in between the ELB and the web server (like a listener or IDP)

QUESTION 42
An enterprise wants to use a third-party SaaS application. The SaaS application needs to have access to issue several API commands to discover Amazon EC2 resources running within the enterprise's account. The enterprise has internal security policies that require any outside access to their environment must conform to the principles of least privilege, and there must be controls in place to ensure that the credentials used by the SaaS vendor cannot be used by any other

third party.
Which of the following would meet all of these conditions:

A.	Create an IAM role for cross-account access, allow the SaaS provider's account to assume the role, and assign it a policy that allows only the actions required by the SaaS application.
B.	From the AWS Management Console navigate to the Security Credentials page and retrieve the access and secret key for your account.
C.	Create an IAM role for EC2 instances, assign it a policy that allows only the actions required for the SaaS application to work, provide the role ARN to the SaaS provider to use when launching their application instances.
D.	Create an IAM user within the enterprise account, assign a user policy to the IAM user that allows only the actions required by the SaaS application, create a new access and secret key for the user and provide these credentials to the SaaS provider.

Answer: A
Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

QUESTION 43
Your company policies require encryption of sensitive data at rest.
You are considering the possible options for protecting data while storing it at rest on an EBS data volume, attached to an EC2 instance.
Which of these options would allow you to encrypt your data at rest? Choose 3 answers

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

Answer: ACD

QUESTION 44
You are designing a photo-sharing mobile app. The application will store all pictures in a single Amazon S3 bucket.
Users will upload pictures from their mobile device directly to Amazon S3 and will be able to view and download their own pictures directly from Amazon S3.
You want to configure security to handle potentially millions of users in the most secure manner possible. What should your server-side application do when a new user registers on the photo-sharing mobile application?

A.	Create an IAM user. Update the bucket policy with appropriate permissions for the IAM user. Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.
B.	Create an IAM user. Assign appropriate permissions to the IAM user. Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.
C.	Create a set of long-term credentials using AWS Security Token Service with appropriate permissions. Store these credentials in the mobile app and use them to access Amazon S3.
D.	Record the user's information in Amazon RDS and create a role in IAM with appropriate permissions.

	When the user uses their mobile app, create temporary credentials using the AWS Security Token Service "AssumeRole" function. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.
E.	Record the user's information in Amazon DynamoDB. When the user uses their mobile app, create temporary credentials using AWS Security Token Service with appropriate permissions. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.

Answer: D
Explanation:
We can use either RDS or DynamoDB, however in our given answers, IAM role is mentioned only with RDS, so I would go with Answer B. Question was explicitly focused on security, so IAM with RDS is the best choice.

QUESTION 45
You are designing a social media site and are considering how to mitigate distributed denial-of-service (DDoS) attacks.
Which of the below are viable mitigation techniques? Choose 3 answers

A.	Use Dedicated Instances to ensure that each Instance has the maximum performance possible.
B.	Add alerts to Amazon CloudWatch to look for high Network In and CPU utilization.
C.	Create processes and capabilities to quickly add and remove rules to the instance OS firewall.
D.	Use an Elastic Load Balancer with auto scaling groups at the web, app, and Amazon Relational Database Service (RDS) tiers.
E.	Use an Amazon CloudFront distribution for both static and dynamic content.
F.	Add multiple elastic network Interfaces (ENIs) to each EC2 instance to Increase the network bandwidth.

Answer: BDE

QUESTION 46
You currently operate a web application in the AWS US-East region. The application runs on an auto- scaled 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. Configure SNS to send log file delivery notifications to your management system. Use IAM roles and S3 bucket policies on the S3 bucket that stores your logs.
B.	Create a new CloudTrail trail with an existing S3 bucket to store the logs and with the global services option selected. Use S3 ACLs and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs.
C.	Create a new CloudTrail trail with one new S3 bucket to store the logs and with the global services option selected. Use IAM roles, S3 bucket policies, and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs.
D.	Create three new CloudTrail trails with three new S3 buckets to store the logs: one for the AWS

	Management Console, one for AWS SDKs, and one for command line tools. Use 1AM roles and S3 bucket policies on the S3 buckets that store your logs.

Answer: C

QUESTION 47
You have an application running on an EC2 instance which will allow users to download files from a private S3 bucket using a pre-signed 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 role for EC2 that allows list access to objects In the S3 bucket; launch the Instance with the role, and retrieve the role's credentials from the EC2 instance metadata.
C.	Create an IAM user for the application with permissions that allow list access to the S3 bucket; the application retrieves the IAM user credentials from a temporary directory with permissions that allow read access only to the Application user.
D.	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.

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

QUESTION 48
A benefits enrollment company is hosting a 3-tier web application running in a VPC on AWS which includes a NAT (Network Address Translation) instance in the public Web tier. There is enough provisioned capacity for the expected workload for the new fiscal year benefit enrollment period plus some extra overhead. Enrollment proceeds nicely for a few days and then the web tier becomes unresponsive. Upon investigation using CloudWatch and other monitoring tools it is discovered that there is an extremely large and unanticipated amount of inbound traffic coming from a set of 15 specific IP addresses over port 80 from a country where the benefits company has no customers. The web tier instances are so overloaded that benefit enrollment administrators cannot even SSH into them.
Which activity would be useful in defending against this attack?

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

Answer: D
Explanation:
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Appendix_NACLs.html

QUESTION 49

You are designing an SSL/TLS solution that requires HTTPS clients to be authenticated by the Web server using client certificate authentication. The solution must be resilient.
Which of the following options would you consider for configuring the Web server infrastructure? Choose 2 answers

A.	Configure your Web servers as the origins for a CloudFront distribution. Use custom SSL certificates on your CloudFront distribution.
B.	Configure ELB with TCP listeners on TCP/443, and place the Web servers behind it.
C.	Configure your Web servers with EIPs. Place the Web servers in a Route53 Record Set, and configure health checks against all Web servers.
D.	Configure ELB with HTTPS listeners, and place the Web servers behind it.

Answer: AD
Explanation:
TCP/443 or HTTPS listener either way you can configure, but you can only upload ssl certificate on HTTPS listener.

QUESTION 50
You've been hired to enhance the overall security posture for a very large e-commerce site. They have a well architected, multi-tier application running in a VPC that uses ELBs in front of both the web and the app tier with static assets served directly from S3. They are using a combination of RDS and DynamoDB 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 1G DirectConnect connection to their VPC. They would then establish Internet connectivity into their space, filter the traffic in a 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 AutoScaling group of EC2 Instances running a host-based WAF. They would redirect Route 53 to resolve to the new WAF tier ELB. The WAF tier would then 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

QUESTION 51
Your company has recently extended its datacenter into a VPC on AWS 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 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.
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-compllant identity provider (IdP) to retrieve temporary security credentials to enable NOC members to sign in to the AWS Management Console.
D.	Use OAuth 2.0 to retrieve temporary AWS security credentials to enable your NOC members to sign in to the AWS Management Console.

Answer: A
Explanation:
http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_enable-console-saml.html

QUESTION 52
Your fortune 500 company has under taken a TCO analysis evaluating the use of Amazon S3 versus acquiring more hardware. The outcome was that all employees would be granted access to use Amazon S3 for storage of their personal documents.
Which of the following will you need to consider so you can set up a solution that incorporates single sign-on from your corporate AD or LDAP directory and restricts access for each user to a designated user folder in a bucket? Choose 3 answers

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

Answer: ADE

QUESTION 53
An administrator is using Amazon CloudFormation to deploy a three tier web application that consists of a web tier and application tier that will utilize Amazon DynamoDB for storage.
When creating the CloudFormation template which of the following would allow the application Instance access to the DynamoDB tables without exposing API credentials?

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

Answer: B

QUESTION 54

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/

QUESTION 55
An AWS customer is deploying an application that is composed of an AutoScaling group of EC2 instances.
The customers security policy requires that every outbound connection from these instances to any other service within the customers Virtual Private Cloud must be authenticated using a unique X.509 certificate that contains the specific Instance-id.
In addition, all X.509 certificates must be signed by the customer's key management service in order to be trusted for authentication.
Which of the following configurations will support these requirements:

A.	Configure an IAM Role that grants access to an Amazon S3 object containing a signed certificate and configure the Auto Scaling group to launch instances with this role. Have the instances bootstrap get the certificate from Amazon S3 upon first boot.
B.	Configure the Auto Scaling group to send an SNS notification of the launch of a new instance to the trusted key management service. Have the key management service generate a signed certificate and send it directly to the newly launched instance.
C.	Embed a certificate into the Amazon Machine Image that is used by the Auto Scaling group. Have the launched instances generate a certificate signature request with the Instance's assigned instance-id to the key management service for signature.
D.	Configure the launched instances to generate a new certificate upon first boot. Have the key management service poll the AutoScaling group for associated instances and send new instances a certificate signature that contains the specific Instance-id.

Answer: A

QUESTION 56
You are looking to migrate your Development (Dev) and Test environments to AWS. You have decided to use separate AWS accounts to host each environment. You plan to link each account's bill to a Master AWS account using Consolidated Billing. To make sure you keep within budget you would like to implement a way for administrators in the Master account to have access to stop, delete and/or terminate resources in both the Dev and Test accounts. Identify which option will allow you to achieve this goal.

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

D.	Create IAM users in the Master account. Create cross-account roles in the Dev and Test accounts that have full Admin permissions and grant the Master account access.

Answer: D
Explanation:
http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_create_for-user.html

QUESTION 57
You are migrating a legacy client-server application to AWS. The application responds to a specific DNS domain (e.g. www.example.com) and has a 2-tier architecture, with multiple application servers and a database server. Remote clients use TCP to connect to the application servers. The application servers need to know the IP address of the clients in order to function properly and are currently taking that information from the TCP socket. A Multi-AZ RDS MySQL instance will be used for the database. During the migration you can change the application code, but you have to file a change request.
How would you implement the architecture on AWS in order to maximize scalability and high availability?

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

Answer: D
Explanation:
https://aws.amazon.com/blogs/aws/elastic-load-balancing-adds-support-for-proxy-protocol/

QUESTION 58
Your company produces customer commissioned one-of-a-kind skiing helmets, combining high fashion with custom technical enhancements. Customers can show off their individuality on the ski slopes and have access to head-up-displays, GPS, rear-view cams and any other technical Innovation they wish to embed in the helmet.
The current manufacturing process is data rich and complex, including assessments to ensure that the custom electronics and materials used to assemble the helmets are to the highest standards. Assessments are a mixture of human and automated assessments.
You need to add a new set of assessment to model the failure modes of the custom electronics using GPUs with CUDA, across a cluster of servers with low latency networking.
What architecture would allow you to automate the existing process using a hybrid approach, and ensure that the architecture can support the evolution of processes over time.

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

C.	Use AWS Data Pipeline to manage movement of data & meta-data and assessments. Use auto-scaling group of C3 with SR-IOV (Single Root I/O Visualization).
D.	Use AWS Data Pipeline to manage movement of data & meta-data and assessments. Use an auto- scaling group of G2 instances in a placement group.

Answer: A

QUESTION 59
A newspaper organization has a 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.	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
B.	Use a single-AZ RDS MySQL instance to store the search index and the JPEG Images, use an EC2 Instance to serve the website and translate user queries into SQL
C.	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 for the website on EC2 instances and use Route53 with DNS round-robin
D.	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
E.	Use S3 with reduced redundancy to store and serve the scanned files, install the commercial search application on EC2 instances and configure with auto-scaling and an Elastic Load Balancer

Answer: D
Explanation:
Cloud search is the perfect option for the search related content.

QUESTION 60
You are designing Internet connectivity for your VPC. The Web servers must be available on the Internet. The application must have a highly available architecture.
Which alternatives should you consider? Choose 2 answers

A.	Assign EIPs to all Web servers. Configure a Route53 record set with all EIPs, with health checks and DNS failover.
B.	Configure a NAT instance in your VPC. Create a default route via the NAT Instance and associate it with all subnets. Configure a DNS A record that points to the NAT Instance public IP address.
C.	Configure a CloudFront distribution and configure the origin to point to the private IP addresses of your Web servers. Configure a Route53 CNAME record to your CloudFront distribution.
D.	Place all your Web servers behind ELB. Configure a Route53 CNAME to point to the ELB DNS name.
E.	Configure ELB with an EIP. Place all your Web servers behind ELB. Configure a Route53 A record that points to the EIP.

Answer: AD

QUESTION 61
A customer has established an AWS Direct Connect connection to AWS. The link is up and routes are being advertised from the customer's end; however the customer is unable to connect from EC2 instances inside its VPC to servers residing in its datacenter.
Which of the following options provide a viable solution to remedy this situation? Choose 2 answers

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

Answer: AD
Explanation:
https://myawsscribble.wordpress.com/2015/09/25/setting-up-and-configuring-aws-directconnect/

QUESTION 62
You have deployed a three-tier web application in a VPC with a CIDR block of 10.0.0.0/28. You initially deploy two web servers, two application servers, two database servers and one NAT instance for a total of seven EC2 instances. The web, application and database servers are deployed across two availability zones (AZs).
You also deploy an ELB in front of the two web servers, and use Route53 for DNS. Web traffic gradually increases in the first few days following the deployment, so you attempt to double the number of instances in each tier of the application to handle the new load.
Unfortunately some of these new Instances fall to launch.
Which of the following could be the root cause? Choose 2 answers

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

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

QUESTION 63
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 this new

connection. After configuring DirectConnect settings in the AWS Console, which of the following options will provide the most seamless transition for your users?

A.	Configure your DirectConnect router, update your VPC route tables to point to the DirectConnect connection, configure your VPN connection with a higher BGP priority, and verify network traffic is leveraging the DirectConnect connection.
B.	Delete your existing VPN connection to avoid routing loops, configure your DirectConnect router with the appropriate settings, and verify network traffic is leveraging DirectConnect.
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 with a higher BGP priority than your VPN router, verify network traffic is leveraging DirectConnect, and then delete your existing VPN connection.

Answer: C
Explanation:
Direct Connect takes priority over Dynamically configured VPN connections.

QUESTION 64
You are implementing AWS Direct Connect.
You intend to use AWS public service endpoints, such as Amazon S3, across the AWS Direct Connect link. You want other Internet traffic to use your existing link to an Internet Service Provider.
What is the correct way to configure AWS Direct Connect for access to services such as Amazon S3?

A.	Create a public interface on your AWS Direct Connect link. Redistribute BGP routes into your existing routing infrastructure; advertise specific routes for your network to AWS.
B.	Create a private interface on your AWS Direct Connect link. Redistribute BGP routes into your existing routing infrastructure and advertise a default route to AWS.
C.	Create a private interface on your AWS Direct Connect link. Configure a static route via your AWS Direct Connect link that points to Amazon S3. Configure specific routes to your network in your VPC.
D.	Configure a public interface on your AWS Direct Connect link. Configure a static route via your AWS Direct Connect link that points to Amazon S3. Advertise a default route to AWS using BGP.

Answer: A
Explanation:
https://aws.amazon.com/directconnect/faqs/

QUESTION 65
What is the name of licensing model in which I can use your existing Oracle Database licenses to run Oracle deployments on Amazon RDS?

A.	Bring Your Own License
B.	Role Bases License
C.	Enterprise License
D.	License Included

Answer: A
Explanation:
https://aws.amazon.com/oracle/

QUESTION 66
When you resize the Amazon RDS DB instance, Amazon RDS will perform the upgrade during the next maintenance window. If you want the upgrade to be performed now, rather than waiting for the maintenance window, specify the _____ option.

A.	ApplyNow
B.	ApplySoon
C.	ApplyThis
D.	ApplyImmediately

Answer: D
Explanation:
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.DBInstance.Modifying.html

QUESTION 67
If I write the below command, what does it do?

```
ec2-run ami-e3a5408a -n 20 -g appserver
```

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

Answer: A

QUESTION 68
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.	Configure Amazon CloudTrail to receive custom logs, use EMR to apply heuristics the logs
B.	Send all the log events to Amazon SQS, setup an Auto Scaling group of EC2 servers to consume the logs and apply the heuristics
C.	Setup an Auto Scaling group of EC2 syslogd servers, store the logs on S3, use EMR to apply heuristics on the logs
D.	Send all the log events to Amazon Kinesis, develop a client process to apply heuristics on the logs

Answer: D
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/

QUESTION 69
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.xlarge 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 and is therefore only done once per week.
Recently, a new chat feature has been implemented in node.js and waits 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 OpsWorks as an application life cycle tool to simplify management of the application and reduce the deployment cycles.
What configuration in AWS OpsWorks is necessary to integrate the new chat module in the most cost-efficient and flexible way?

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

Answer: C

QUESTION 70
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 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.
B.	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.
C.	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.
D.	Create your RDS instance separately and pass its DNS name to your app's DB connection string as an environment variable. 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.

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.
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/AWSHowTo.RDS.html

QUESTION 71
Your company has an on-premises, multi-tier PHP web application, which recently experienced downtime due to a large burst in web traffic due to a company announcement. Over the coming days, you're expecting similar announcements to drive similar unpredictable bursts, and are looking to find ways to quickly improve your infrastructures ability to handle unexpected increases in traffic. The application currently consists of 2 tiers: A web tier, which consists of a load balancer and several Linux Apache web servers, as well as a database tier, which hosts a Linux server hosting a MySQL database.
Which scenario below will provide full site functionality, while helping to improve the availability of your application in the short timeframe required?

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

Answer: C
Explanation:
You can have CloudFront sit in front of your on-prem web environment, via a custom origin (the origin doesn't have to be in AWS). This would protect against unexpected bursts in traffic by letting CloudFront handle the traffic that it can out of cache, thus hopefully removing some of the load from your on-prem web servers.

QUESTION 72
A read only news reporting site with a combined web and application tier and a database tier that receives large and unpredictable traffic demands must be able to respond to these traffic fluctuations automatically. What AWS services should be used meet these requirements?

A.	Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaling group monitored with CloudWatch, and RDS with read replicas
B.	Stateful instances for the web and application tier in an autoscaling group monitored with CloudWatch, and multi-AZ RDS
C.	Stateful instances for the web and application tier in an autoscaling group monitored with CloudWatch, and RDS with read replicas
D.	Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaling group monitored with CloudWatch, and multi-AZ RDS

Answer: A
Explanation:

"A readonly reporting site" - so stateless and read-replicas can be used to scale. Multi-AZ will not provide the scaling requirements.

QUESTION 73
Your company is getting ready to do a major public announcement of a social media site on AWS. The website is running on EC2 instances deployed across multiple Availability Zones with an Multi-AZ RDS MySQL Extra Large DB Instance backend. The site performs a high number of small reads and writes per second and relies on an eventual consistency model. After comprehensive tests you discover that there is read contention on RDS MySQL.
Which are the best approaches to meet these requirements? Choose 2 answers

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

Answer: AB

QUESTION 74
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 CloudFront 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 Restrict Viewer Access, Forward Query String set to true and minimum TTL of 0.
D.	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.

Answer: D

QUESTION 75
You require the ability to analyze a customer's clickstream data on a website, so they can do behavioral analysis. Your customer needs to know what sequence of pages and ads their customer clicked on. This data will be used in real time to modify the page layouts as customers dick through the site, to increase stickiness and advertising click-through.
Which option meets the requirements for capturing and analyzing this data?

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

Answer: C
Explanation:
http://www.slideshare.net/AmazonWebServices/aws-webcast-introduction-to-amazon-kinesis

QUESTION 76
A large real-estate brokerage is exploring the option of 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 offers 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 send device location using SQS, EC2 instances will retrieve the relevant offers from DynamoDB. AWS Mobile Push will be used to send offers to the mobile application.
B.	Use AWS DirectConnect or VPN to establish connectivity with mobile carriers. EC2 instances will receive the mobile applications location through earner connection; RDS will be used to store and retrieve relevant offers. EC2 instances will communicate with mobile carriers to push alerts back to the mobile application.
C.	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 offers. EC2 instances will communicate with mobile carriers/device providers to push alerts back to mobile application.
D.	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 mobile carriers/device providers to push alerts back to the mobile application.

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

QUESTION 77
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 text 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 files, SQS to distribute elaboration commands to a group of hosts working in parallel, Auto Scaling to dynamically size the group of hosts depending on the length of the SQS queue.
B.	S3 to store I/O files, 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.
C.	EBS with Provisioned IOPS (PIOPS) to store I/O files, 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.

D.	EBS with Provisioned IOPS (PIOPS) to store I/O files, SQS to distribute elaboration commands to a group of hosts working in parallel. Auto Scaling to dynamically size the group of hosts depending on the length of the SQS queue.

Answer: D

QUESTION 78
You are implementing a URL whitelisting system for a company that wants to restrict outbound HTTP/S connections to specific domains from their EC2-hosted applications.
You deploy a single EC2 instance running proxy software and configure it to accept traffic from all subnets and EC2 instances in the VPC.
You configure the proxy to only pass through traffic to domains that you define in its whitelist configuration.
You have a nightly maintenance window of 10 minutes where all instances fetch new software updates. Each update is about 200MB in size and there are 500 instances in the VPC that routinely fetch updates. After a few days you notice that some machines are falling to successfully download some, but not all, of their updates within the maintenance window.
The download URLs used for these updates are correctly listed in the proxy's whitelist configuration and you are able to access them manually using a web browser on the instances.
What might be happening? Choose 2 answers

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

Answer: AD
Explanation:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-ec2-config.html

QUESTION 79
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.

QUESTION 80
To serve Web traffic for a popular product, your chief financial officer and IT director have purchased 10 m1.large heavy utilization Reserved Instances (RIs), evenly spread across two availability zones; Route 53 is used to deliver the traffic to an Elastic Load Balancer (ELB). After several months, the product grows even more popular and you need additional capacity. As a result, your company purchases two c3.2xlarge medium utilization RIs.
You register the two c3.2xlarge instances with your ELB and quickly find that the m1.large instances are at 100% of capacity and the c3.2xlarge instances have significant capacity that's unused. Which option is the most cost effective and uses EC2 capacity most effectively?

A.	Configure Autoscaling group and Launch Configuration with ELB to add up to 10 more on-demand m1.large instances when triggered by Cloudwatch. Shut off c3.2xlarge instances.
B.	Configure ELB with two c3.2xlarge instances and use on-demand Autoscaling group for up to two additional c3.2xlarge instances. Shut off m1.large instances.
C.	Route traffic to EC2 m1.large and c3.2xlarge instances directly using Route 53 latency based routing and health checks. Shut off ELB.
D.	Use a separate ELB for each instance type and distribute load to ELBs with Route 53 weighted round robin.

Answer: D
Explanation:
http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html

QUESTION 81
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 3 answers

A.	Configure Amazon CloudFront dynamic content support to enable caching of re-usable content from your site.
B.	Set up a second installation in another region, and use the Amazon Route 53 latency-based routing feature to select the right region.

C.	Lower the scale up trigger of your Auto Scaling group to 30% so it scales more aggressively.
D.	Add an Amazon ElastiCache caching layer to your application for storing sessions and frequent DB queries.
E.	Switch the Amazon RDS database to the high memory extra large instance type.

Answer: ADE

QUESTION 82
You are responsible for a legacy web application whose server environment is approaching end of life.
You would like to migrate this application to AWS as quickly as possible, since the application environment currently has the following limitations:

```
- the VM's single 10GB VMDK is almost full;
- the virtual network Interface still uses the 10Mbps dnver, which
leaves your 100Mbps WAN connection completely underutilized;
- it is currently running on a highly customized, Windows VM within a
VMware environment; ?you do not have the installation media.
```

This is a mission critical application with an RTO (Recovery Time Objective) of 8 hours, RPO (Recovery Point Objective) of 1 hour.
How could you best migrate this application to AWS while meeting your business continuity requirements?

A.	Use S3 to create a backup of the VM and restore the data into EC2.
B.	Use the EC2 VM Import Connector for vCenter to import the VM into EC2.
C.	Use the ec2-bundle-instance API to import an image of the VM into EC2.
D.	Use Import/Export to import the VM as an EBS snapshot and attach to EC2.

Answer: B
Explanation:
https://aws.amazon.com/developertools/2759763385083070

QUESTION 83
Your company hosts a social media site supporting users in multiple countries.
You have been asked to provide a highly available design for the application that leverages multiple regions for the most recently accessed content and latency sensitive portions of the web 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.	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 ElastiCache cluster in each region and leverage SNS notifications to propagate user preference changes to a worker node in each region.
B.	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 DB to each ElastiCache cluster.
C.	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 SQS workers for propagating updates to each table.
D.	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 SQS workers for propagating DynamoDB updates.

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

QUESTION 84
A web company is looking to implement an external payment service into their highly available application deployed in a VPC. Their application EC2 instances are behind a public facing ELB. Auto Scaling is used to add additional instances as traffic Increases. Under normal load the application runs 2 Instances in the Auto Scaling group but at peak it can scale 3x in size.
The application instances need to communicate with the payment service over the Internet, which requires whitelisting of all public IP addresses used to communicate with it. A maximum of 4 whitelisted IP addresses are allowed at a time and can be added through an API.
How should they architect their solution?

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

Answer: C

QUESTION 85
A company is building a voting system for a popular TV show, viewers will watch the performances then visit the show's website to vote for their favorite performer. It is expected that in a short period of time after the show has finished the site will receive millions of visitors, the visitors will first login to the site using theirAmazon.com credentials and then submit their vote. After the voting is completed the page will display the vote totals. The company needs to build the site such that it can handle the rapid influx of traffic while maintaining good performance but also wants to keep costs to a minimum.
Which of the design patters below should they use?

A.	Use CloudFront and an Elastic Load Balancer in front of an auto-scaled set of web servers, the web servers will first call the Login With Amazon service to authenticate the user, the web servers will process the users vote and store the result into a DynamoDB table using IAM Roles for EC2 Instances to gain permissions to the DynamoDB table.
B.	Use CloudFront and an Elastic Load Balancer in front of an auto-scaled set of web servers, the web servers will first call the Login With Amazon service to authenticate the user, the web servers will process the users vote and store the result into an SQS queue using IAM Roles for EC2 Instances to gain permissions to the SQS queue. A set of application servers will then retrieve the items from the queue and store the result into a

	DynamoDB table.
C.	Use CloudFront and an Elastic Load Balancer in front of an auto-scaled set of web servers, the web servers will first call the Login With Amazon service to authenticate the user then process the users vote and store the result into a multi-AZ Relational Database Service instance.
D.	Use CloudFront and the static website hosting feature of S3 with the Javascript SDK to call the Login with Amazon service to authenticate the user, use IAM Roles to gain permissions to a DynamoDB table to store the users vote.

Answer: B

QUESTION 86
You are developing a new mobile application and are considering storing user preferences in AWS. This would provide a more uniform cross-device experience to users using multiple mobile devices to access the application. The preference data for each user is estimated to be 50KB in size. Additionally, 5 million customers are expected to use the application on a regular basis.
The solution needs to be cost- effective, highly-available, scalable and secure.
How would you design a solution to meet the above requirements?

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

Answer: D
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.

QUESTION 87
The _____ service is targeted at organizations with multiple users or systems that use AWS products such as Amazon EC2, Amazon SimpleDB, and the AWS Management Console.

A.	Amazon RDS

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

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

QUESTION 88
Select the correct set of options.
These are the initial settings for the default security group:

A.	Allow no inbound traffic, Allow all outbound traffic and Allow instances associated with this security group to talk to each other
B.	Allow all inbound traffic, Allow no outbound traffic and Allow instances associated with this security group to talk to each other
C.	Allow no inbound traffic, Allow all outbound traffic and Does NOT allow instances associated with this security group to talk to each other
D.	Allow all inbound traffic, Allow all outbound traffic and Does NOT allow instances associated with this security group to talk to each other

Answer: A
Explanation:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html#default-security-group
A default security group is named default, and it has an ID assigned by AWS. The following are the initial settings for each default security group:
Allow inbound traffic only from other instances associated with the default security group
Allow all outbound traffic from the instance
The default security group specifies itself as a source security group in its inbound rules. This is what allows instances associated with the default security group to communicate with other instances associated with the default security group.

QUESTION 89
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.	Both A and B
B.	None of these
C.	VPC Addresses
D.	EC2 Addresses

Answer: D
Explanation:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html
Stopping an instance
EC2-Classic
If you stop an instance, its Elastic IP address is disassociated, and you must reassociate the Elastic IP address when you restart the instance.
EC2-VPC
If you stop an instance, its Elastic IP address remains associated.

QUESTION 90
You are designing a data leak prevention solution for your VPC environment.
You want your VPC instances to be able to access software depots and distributions on the Internet for product updates. The depots and distributions are accessible via third party CDNs by their URLs.
You want to explicitly deny any other outbound connections from your VPC instances to hosts on the Internet.
Which of the following options would you consider?

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

Answer: B
Explanation:
Organizations usually implement proxy solutions to provide URL and web content filtering, IDS/IPS, data loss prevention, monitoring, and advanced threat protection.
https://d0.awsstatic.com/aws-answers/Controlling_VPC_Egress_Traffic.pdf

QUESTION 91
By default, Amazon Cognito maintains the last-written version of the data.
You can override this behavior and resolve data conflicts programmatically.
In addition, push synchronization allows you to use Amazon Cognito to send a silent _____ notification to all devices associated with an identity to notify them that new data is available.

A.	get
B.	post
C.	pull
D.	push

Answer: D
Explanation:
http://aws.amazon.com/cognito/faqs/

QUESTION 92
You want to use AWS CodeDeploy to deploy an application to Amazon EC2 instances running within an Amazon Virtual Private Cloud (VPC).
What criterion must be met for this to be possible?

A.	The AWS CodeDeploy agent installed on the Amazon EC2 instances must be able to access only the public AWS CodeDeploy endpoint.
B.	The AWS CodeDeploy agent installed on the Amazon EC2 instances must be able to access only the public Amazon S3 service endpoint.
C.	The AWS CodeDeploy agent installed on the Amazon EC2 instances must be able to access the public AWS CodeDeploy and Amazon S3 service endpoints.
D.	It is not currently possible to use AWS CodeDeploy to deploy an application to Amazon EC2 instances running within an Amazon Virtual Private Cloud (VPC.)

Answer: C
Explanation:
You can use AWS CodeDeploy to deploy an application to Amazon EC2 instances running within an Amazon Virtual Private Cloud (VPC).
However, the AWS CodeDeploy agent installed on the Amazon EC2 instances must be able to access the public AWS CodeDeploy and Amazon S3 service endpoints.
http://aws.amazon.com/codedeploy/faqs/

QUESTION 93
An IAM user is trying to perform an action on an object belonging to some other root account's bucket.
Which of the below mentioned options will AWS S3 not verify?

A.	The object owner has provided access to the IAM user
B.	Permission provided by the parent of the IAM user on the bucket
C.	Permission provided by the bucket owner to the IAM user
D.	Permission provided by the parent of the IAM user

Answer: B
Explanation:
If the IAM user is trying to perform some action on the object belonging to another AWS user's bucket, S3 will verify whether the owner of the IAM user has given sufficient permission to him. It also verifies the policy for the bucket as well as the policy defined by the object owner.
http://docs.aws.amazon.com/AmazonS3/latest/dev/access-control-auth-workflow-object-operation.html

QUESTION 94
An organization is planning to extend their data center by connecting their DC with the AWS VPC using the VPN gateway. The organization is setting up a dynamically routed VPN connection.
Which of the below mentioned answers is not required to setup this configuration?

A.	The type of customer gateway, such as Cisco ASA, Juniper J-Series, Juniper SSG, Yamaha.
B.	Elastic IP ranges that the organization wants to advertise over the VPN connection to the VPC.
C.	Internet-routable IP address (static) of the customer gateway's external interface.
D.	Border Gateway Protocol (BGP) Autonomous System Number (ASN) of the customer gateway.

Answer: B
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. The organization wants to extend their network into the cloud and also directly access the internet from their AWS VPC. Thus, the organization should setup a Virtual Private Cloud (VPC) with a public subnet and a private subnet, and a virtual private gateway to enable communication with their data center network over an IPsec VPN tunnel. To setup this configuration the organization needs to use the Amazon VPC with a VPN connection. The organization network administrator must designate a physical appliance as a customer gateway and configure it. The organization would need the below mentioned information to setup this configuration:
The type of customer gateway, such as Cisco ASA, Juniper J-Series, Juniper SSG, Yamaha Internet-routable IP address (static) of the customer gateway's external interface Border Gateway Protocol (BGP) Autonomous System Number (ASN) of the customer gateway, if the organization is creating a dynamically routed VPN connection. Internal network IP ranges that the user wants to advertise over the VPN connection to the VPC.

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

QUESTION 95
In the context of AWS IAM, identify a true statement about user passwords (login profiles).

A.	They must contain Unicode characters.
B.	They can contain any Basic Latin (ASCII) characters.
C.	They must begin and end with a forward slash (/).
D.	They cannot contain Basic Latin (ASCII) characters.

Answer: B
Explanation:
The user passwords (login profiles) of IAM users can contain any Basic Latin (ASCII) characters.
http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html

QUESTION 96
An organization is planning to host a Wordpress blog as well a joomla CMS on a single instance launched with VPC. The organization wants to have separate domains for each application and assign them using Route 53. The organization may have about ten instances each with two applications as mentioned above. While launching the instance, the organization configured two separate network interfaces (primary + ENI) and wanted to have two elastic IPs for that instance. It was suggested to use a public IP from AWS instead of an elastic IP as the number of elastic IPs is restricted. What action will you recommend to the organization?

A.	I agree with the suggestion but will prefer that the organization should use separate subnets with each ENI for different public IPs.
B.	I do not agree as it is required to have only an elastic IP since an instance has more than one ENI and AWS does not assign a public IP to an instance with multiple ENIs.
C.	I do not agree as AWS VPC does not attach a public IP to an ENI; so the user has to use only an elastic IP only.
D.	I agree with the suggestion and it is recommended to use a public IP from AWS since the organization is going to use DNS with Route 53.

Answer: B
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. An Elastic Network Interface (ENI) is a virtual network interface that the user can attach to an instance in a VPC. The user can attach up to two ENIs with a single instance. However, AWS cannot assign a public IP when there are two ENIs attached to a single instance. It is recommended to assign an elastic IP in this scenario. If the organization wants more than 5 EIPs they can request AWS to increase the number.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 97
What is the default maximum number of VPCs allowed per region?

A.	5
B.	10
C.	100
D.	15

Answer: A
Explanation:
The maximum number of VPCs allowed per region is 5.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Appendix_Limits.html

QUESTION 98
A customer has a website which shows all the deals available across the market. The site experiences a load of 5 large EC2 instances generally.
However, a week before Thanksgiving vacation they encounter a load of almost 20 large instances. The load during that period varies over the day based on the office timings.
Which of the below mentioned solutions is cost effective as well as help the website achieve better performance?

A.	Setup to run 10 instances during the pre-vacation period and only scale up during the office time by launching 10 more instances using the AutoScaling schedule.
B.	Keep only 10 instances running and manually launch 10 instances every day during office hours.
C.	During the pre-vacation period setup 20 instances to run continuously.
D.	During the pre-vacation period setup a scenario where the organization has 15 instances running and 5 instances to scale up and down using Auto Scaling based on the network I/O policy.

Answer: D
Explanation:
AWS provides an on demand, scalable infrastructure. AWS EC2 allows the user to launch On-Demand instances and the organization should create an AMI of the running instance. When the organization is experiencing varying loads and the time of the load is not known but it is higher than the routine traffic it is recommended that the organization launches a few instances before hand and then setups AutoScaling with policies which scale up and down as per the EC2 metrics, such as Network I/O or CPU utilization. If the organization keeps all 10 additional instances as a part of the AutoScaling policy sometimes during a sudden higher load it may take time to launch instances and may not give an optimal performance. This is the reason it is recommended that the organization keeps an additional 5 instances running and the next 5 instances scheduled as per the AutoScaling policy for cost effectiveness.
http://media.amazonwebservices.com/AWS_Web_Hosting_Best_Practices.pdf

QUESTION 99
An organization is setting a website on the AWS VPC. The organization has blocked a few IPs to avoid a D-DOS attack.
How can the organization configure that a request from the above mentioned IPs does not access the application instances?

A.	Create an IAM policy for VPC which has a condition to disallow traffic from that IP address.
B.	Configure a security group at the subnet level which denies traffic from the selected IP.
C.	Configure the security group with the EC2 instance which denies access from that IP address.
D.	Configure an ACL at the subnet which denies the traffic from that IP address.

Answer: D
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. AWS provides two features that the user can use to increase security in VPC: security groups and network ACLs. Security group works at the instance level while ACL works at the subnet level. ACL allows both allow and deny rules. Thus, when the user wants to reject traffic from the selected IPs it is

recommended to use ACL with subnets.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_ACLs.html

QUESTION 100
An organization has 4 people in the IT operations team who are responsible to manage the AWS infrastructure. The organization wants to setup that each user will have access to launch and manage an instance in a zone which the other user cannot modify. Which of the below mentioned options is the best solution to set this up?

A.	Create four AWS accounts and give each user access to a separate account.
B.	Create an IAM user and allow them permission to launch an instance of a different sizes only.
C.	Create four IAM users and four VPCs and allow each IAM user to have access to separate VPCs.
D.	Create a VPC with four subnets and allow access to each subnet for the individual IAM user.

Answer: D
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. The VPC also work with IAM and the organization can create IAM users who have access to various VPC services. The organization can setup access for the IAM user who can modify the security groups of the VPC. The sample policy is given below:
{
"Version": "2012-10-17",
"Statement":
[{ "Effect": "Allow",
"Action": "ec2:RunInstances",
"Resource":
["arn:aws:ec2:region::image/ami-*",
"arn:aws:ec2:region:account:subnet/subnet-1a2b3c4d",
"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/sg-123abc123"] }]
}
With this policy the user can create four subnets in separate zones and provide IAM user access to each subnet.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_IAM.html

QUESTION 101
An organization is planning to host an application on the AWS VPC. The organization wants dedicated instances. However, an AWS consultant advised the organization not to use dedicated instances with VPC as the design has a few limitations. Which of the below mentioned statements is not a limitation of dedicated instances with VPC?

A.	All instances launched with this VPC will always be dedicated instances and the user cannot use a default tenancy model for them.
B.	It does not support the AWS RDS with a dedicated tenancy VPC.
C.	The user cannot use Reserved Instances with a dedicated tenancy model.
D.	The EBS volume will not be on the same tenant hardware as the EC2 instance though the user has configured dedicated tenancy.

Answer: C
Explanation:

The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Dedicated instances are Amazon EC2 instances that run in a Virtual Private Cloud (VPC) on hardware that is dedicated to a single customer. The client's dedicated instances are physically isolated at the host hardware level from instances that are not dedicated instances as well as from instances that belong to other AWS accounts. All instances launched with the dedicated tenancy model of VPC will always be dedicated instances. Dedicated tenancy has a limitation that it may not support a few services, such as RDS. Even the EBS will not be on dedicated hardware. However the user can save some cost as well as reserve some capacity by using a Reserved Instance model with dedicated tenancy.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/dedicated-instance.html

QUESTION 102
In which step of using AWS Direct Connect should the user determine the required port speed?

A.	Complete the Cross Connect
B.	Verify Your Virtual Interface
C.	Download Router Configuration
D.	Submit AWS Direct Connect Connection Request

Answer: D
Explanation:
To submit an AWS Direct Connect connection request, you need to provide the following information:
Your contact information.
The AWS Direct Connect Location to connect to.
Details of AWS Direct Connect partner if you use the AWS Partner Network (APN) service.
The port speed you require, either 1 Gbps or 10 Gbps.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/getstarted.html#ConnectionRequest

QUESTION 103
In Amazon IAM, what is the maximum length for a role name?

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

Answer: C
Explanation:
In Amazon IAM, the maximum length for a role name is 64 characters.
http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html

QUESTION 104
A user is planning to host a web server as well as an app server on a single EC2 instance which is a part of the public subnet of a VPC. How can the user setup to have two separate public IPs and separate security groups for both the application as well as the web server?

A.	Launch VPC with two separate subnets and make the instance a part of both the subnets.
B.	Launch a VPC instance with two network interfaces. Assign a separate security group and elastic

	IP to them.
C.	Launch a VPC instance with two network interfaces. Assign a separate security group to each and AWS will assign a separate public IP to them.
D.	Launch a VPC with ELB such that it redirects requests to separate VPC instances of the public subnet.

Answer: B
Explanation:
If you need to host multiple websites(with different IPs) on a single EC2 instance, the following is the suggested method from AWS.
Launch a VPC instance with two network interfaces
Assign elastic IPs from VPC EIP pool to those interfaces (Because, when the user has attached more than one network interface with an instance, AWS cannot assign public IPs to them.) Assign separate Security Groups if separate Security Groups are needed This scenario also helps for operating network appliances, such as firewalls or load balancers that have multiple private IP addresses for each network interface.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/MultipleIP.html

QUESTION 105
You have subscribed to the AWS Business and Enterprise support plan.
Your business has a backlog of problems, and you need about 20 of your IAM users to open technical support cases.
How many users can open technical support cases under the AWS Business and Enterprise support plan?

A.	5 users
B.	10 users
C.	Unlimited
D.	1 user

Answer: C
Explanation:
In the context of AWS support, the Business and Enterprise support plans allow an unlimited number of users to open technical support cases (supported by AWS Identity and Access Management (IAM)).
https://aws.amazon.com/premiumsupport/faqs/

QUESTION 106
While implementing the policy keys in AWS Direct Connect, if you use _____ and the request comes from an Amazon EC2 instance, the instance's public IP address is evaluated to determine if access is allowed.

A.	aws:SecureTransport
B.	aws:EpochIP
C.	aws:SourceIp
D.	aws:CurrentTime

Answer: C
Explanation:
While implementing the policy keys in Amazon RDS, if you use aws:SourceIp and the request comes from an Amazon EC2 instance, the instance's public IP address is evaluated to determine if access is allowed.

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

QUESTION 107
How many g2.2xlarge on-demand instances can a user run in one region without taking any limit increase approval from AWS?

A.	20
B.	2
C.	5
D.	10

Answer: C
Explanation:
Generally AWS EC2 allows running 20 on-demand instances and 100 spot instances at a time. This limit can be increased by requesting at https://aws.amazon.com/contact-us/ec2-request. Excluding certain types of instances, the limit is lower than mentioned above. For g2.2xlarge, the user can run only 5 on-demand instance at a time.
http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_ec2

QUESTION 108
A user has created a MySQL RDS instance with PIOPS. Which of the below mentioned statements will help user understand the advantage of PIOPS?

A.	The user can achieve additional dedicated capacity for the EBS I/O with an enhanced RDS option
B.	It uses a standard EBS volume with optimized configuration the stacks
C.	It uses optimized EBS volumes and optimized configuration stacks
D.	It provides a dedicated network bandwidth between EBS and RDS

Answer: C
Explanation:
RDS DB instance storage comes in two types: standard and provisioned IOPS. Standard storage is allocated on the Amazon EBS volumes and connected to the user's DB instance. Provisioned IOPS uses optimized EBS volumes and an optimized configuration stack. It provides additional, dedicated capacity for the EBS I/O.
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html

QUESTION 109
A user authenticating with Amazon Cognito will go through a multi-step process to bootstrap their credentials.
Amazon Cognito has two different flows for authentication with public providers. Which of the following are the two flows?

A.	Authenticated and non-authenticated
B.	Public and private
C.	Enhanced and basic
D.	Single step and multistep

Answer: C
Explanation:
A user authenticating with Amazon Cognito will go through a multi-step process to bootstrap their credentials. Amazon Cognito has two different flows for authentication with public providers:

enhanced and basic.
http://docs.aws.amazon.com/cognito/devguide/identity/concepts/authentication-flow/

QUESTION 110
Which of the following is the Amazon Resource Name (ARN) condition operator that can be used within an Identity and Access Management (IAM) policy to check the case-insensitive matching of the ARN?

A.	ArnCheck
B.	ArnMatch
C.	ArnCase
D.	ArnLike

Answer: D
Explanation:
Amazon Resource Name (ARN) condition operators let you construct Condition elements that restrict access based on comparing a key to an ARN. ArnLike, for instance, is a case-insensitive matching of the ARN. Each of the six colon-delimited components of the ARN is checked separately and each can include a multi-character match wildcard (*) or a single-character match wildcard (?).
http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html

QUESTION 111
An organization is creating a VPC for their application hosting. The organization has created two private subnets in the same AZ and created one subnet in a separate zone.
The organization wants to make a HA system with the internal ELB. Which of these statements is true with respect to an internal ELB in this scenario?

A.	ELB can support only one subnet in each availability zone.
B.	ELB does not allow subnet selection; instead it will automatically select all the available subnets of the VPC.
C.	If the user is creating an internal ELB, he should use only private subnets.
D.	ELB can support all the subnets irrespective of their zones.

Answer: A
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud.
The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances.
There are two ELBs available with VPC: internet facing and internal (private) ELB. For internal servers, such as App servers the organization can create an internal load balancer in their VPC and then place back-end application instances behind the internal load balancer.
The internal load balancer will route requests to the back-end application instances, which are also using private IP addresses and only accept requests from the internal load balancer.
The Internal ELB supports only one subnet in each AZ and asks the user to select a subnet while configuring internal ELB.
http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/USVPC_creating_basic_lb.ht ml

QUESTION 112

In Amazon ElastiCache, the failure of a single cache node can have an impact on the availability of your application and the load on your back-end database while ElastiCache provisions a replacement for the failed cache node and it get repopulated. Which of the following is a solution to reduce this potential availability impact?

A.	Spread your memory and compute capacity over fewer number of cache nodes, each with smaller capacity.
B.	Spread your memory and compute capacity over a larger number of cache nodes, each with smaller capacity.
C.	Include fewer number of high capacity nodes.
D.	Include a larger number of cache nodes, each with high capacity.

Answer: B
Explanation:
In Amazon ElastiCache, the number of cache nodes in the cluster is a key factor in the availability of your cluster running Memcached. The failure of a single cache node can have an impact on the availability of your application and the load on your back-end database while ElastiCache provisions a replacement for the failed cache node and it get repopulated.
You can reduce this potential availability impact by spreading your memory and compute capacity over a larger number of cache nodes, each with smaller capacity, rather than using a fewer number of high capacity nodes.
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/CacheNode.Memcached.html

QUESTION 113
MapMySite is setting up a web application in the AWS VPC. The organization has decided to use an AWS RDS instead of using its own DB instance for HA and DR requirements.
The organization also wants to secure RDS access. How should the web application be setup with RDS?

A.	Create a VPC with one public and one private subnet. Launch an application instance in the public subnet while RDS is launched in the private subnet.
B.	Setup a public and two private subnets in different AZs within a VPC and create a subnet group. Launch RDS with that subnet group.
C.	Create a network interface and attach two subnets to it. Attach that network interface with RDS while launching a DB instance.
D.	Create two separate VPCs and launch a Web app in one VPC and RDS in a separate VPC and connect them with VPC peering.

Answer: B
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources, such as RDS into a virtual network that the user has defined. Subnets are segments of a VPC's IP address range that the user can designate to a group of VPC resources based on the security and operational needs.
A DB subnet group is a collection of subnets (generally private) that a user can create in a VPC and assign to the RDS DB instances. A DB subnet group allows the user to specify a particular VPC when creating the DB instances. Each DB subnet group should have subnets in at least two Availability Zones in a given region.
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.html

QUESTION 114
When does an AWS Data Pipeline terminate the AWS Data Pipeline-managed compute resources?

A.	AWS Data Pipeline terminates AWS Data Pipeline-managed compute resources every 2 hours.
B.	When the final activity that uses the resources is running
C.	AWS Data Pipeline terminates AWS Data Pipeline-managed compute resources every 12 hours.
D.	When the final activity that uses the resources has completed successfully or failed

Answer: D
Explanation:
Compute resources will be provisioned by AWS Data Pipeline when the first activity for a scheduled time that uses those resources is ready to run, and those instances will be terminated when the final activity that uses the resources has completed successfully or failed.
https://aws.amazon.com/datapipeline/faqs/

QUESTION 115
What bandwidths do AWS Direct Connect currently support?

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

Answer: D
Explanation:
AWS Direct Connection currently supports 1Gbps and 10 Gbps.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

QUESTION 116
The Principal element of an IAM policy refers to the specific entity that should be allowed or denied permission, whereas the _____ translates to everyone except the specified entity.

A.	NotPrincipal
B.	Vendor
C.	Principal
D.	Action

Answer: A
Explanation:
The element NotPrincipal that is included within your IAM policy statements allows you to specify an exception to a list of principals to whom the access to a specific resource is either allowed or denied. Use the NotPrincipal element to specify an exception to a list of principals. For example, you can deny access to all principals except the one named in the NotPrincipal element.
http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html#Principal

QUESTION 117
Doug has created a VPC with CIDR 10.201.0.0/16 in his AWS account. In this VPC he has created a public subnet with CIDR block 10.201.31.0/24.
While launching a new EC2 from the console, he is not able to assign the private IP address 10.201.31.6 to this instance.
Which is the most likely reason for this issue?

A.	Private address IP 10.201.31.6 is currently assigned to another interface.

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

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

QUESTION 118
A user is configuring MySQL RDS with PIOPS. What should be the minimum size of DB storage provided by the user?

A.	1 TB
B.	50 GB
C.	5 GB
D.	100 GB

Answer: D
Explanation:
If the user is trying to enable PIOPS with MySQL RDS, the minimum size of storage should be 100 GB.
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.html

QUESTION 119
The Statement element, of an AWS IAM policy, contains an array of individual statements. Each individual statement is a(n) _____ block enclosed in braces { }.

A.	XML
B.	JavaScript
C.	JSON
D.	AJAX

Answer: C
Explanation:
The Statement element, of an IAM policy, contains an array of individual statements. Each individual statement is a JSON block enclosed in braces { }.
http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html

QUESTION 120
If no explicit deny is found while applying IAM's Policy Evaluation Logic, the enforcement code looks for any _____ instructions that would apply to the request.

A.	"cancel"
B.	"suspend"
C.	"allow"

D.	"valid"

Answer: C
Explanation:
If an explicit deny is not found among the applicable policies for a specific request, IAM's Policy Evaluation Logic checks for any "allow" instructions to check if the request can be successfully completed.
http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_EvaluationLogic.html

QUESTION 121
An organization is hosting a scalable web application using AWS. The organization has configured ELB and Auto Scaling to make the application scalable.
Which of the below mentioned statements is not required to be followed for ELB when the application is planning to host a web application on VPC?

A.	The ELB and all the instances should be in the same subnet.
B.	Configure the security group rules and network ACLs to allow traffic to be routed between the subnets in the VPC.
C.	The internet facing ELB should have a route table associated with the internet gateway.
D.	The internet facing ELB should be only in a public subnet.

Answer: A
Explanation:
Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances. There are two ELBs available with VPC: internet facing and internal (private) ELB. For the internet facing ELB it is required that the ELB should be in a public subnet. After the user creates the public subnet, he should ensure to associate the route table of the public subnet with the internet gateway to enable the load balancer in the subnet to connect with the internet. The ELB and instances can be in a separate subnet. However, to allow communication between the instance and the ELB the user must configure the security group rules and network ACLs to allow traffic to be routed between the subnets in his VPC.
http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/CreateVPCForELB.html

QUESTION 122
An organization (account ID 123412341234) has configured the IAM policy to allow the user to modify his credentials.
What will the below mentioned statement allow the user to perform?

```
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow",
"Action": [
"iam:AddUserToGroup",
"iam:RemoveUserFromGroup",
"iam:GetGroup"
],
"Resource": "arn:aws:iam:: 123412341234:group/TestingGroup"
}
```

]

A.	Allow the IAM user to update the membership of the group called TestingGroup
B.	The IAM policy will throw an error due to an invalid resource name
C.	The IAM policy will allow the user to subscribe to any IAM group
D.	Allow the IAM user to delete the TestingGroup

Answer: A
Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the organization (account ID 123412341234) wants their users to manage their subscription to the groups, they should create a relevant policy for that. The below mentioned policy allows the respective IAM user to update the membership of the group called MarketingGroup.
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow",
"Action": [
"iam:AddUserToGroup",
"iam:RemoveUserFromGroup",
"iam:GetGroup"
],
"Resource": "arn:aws:iam:: 123412341234:group/ TestingGroup " }]
http://docs.aws.amazon.com/IAM/latest/UserGuide/Credentials-Permissions-examples.html#creds-polici es-credentials

QUESTION 123
A user has configured EBS volume with PIOPS. The user is not experiencing the optimal throughput.
Which of the following could not be factor affecting I/O performance of that EBS volume?

A.	EBS bandwidth of dedicated instance exceeding the PIOPS
B.	EBS volume size
C.	EC2 bandwidth
D.	Instance type is not EBS optimized

Answer: B
Explanation:
If the user is not experiencing the expected IOPS or throughput that is provisioned, ensure that the EC2 bandwidth is not the limiting factor, the instance is EBS-optimized (or include 10 Gigabit network connectivity) and the instance type EBS dedicated bandwidth exceeds the IOPS more than he has provisioned.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-io-characteristics.html

QUESTION 124
How can multiple compute resources be used on the same pipeline in AWS Data Pipeline?

A.	You can use multiple compute resources on the same pipeline by defining multiple cluster objects in your definition file and associating the cluster to use for each activity via its runsOn field.
B.	You can use multiple compute resources on the same pipeline by defining multiple cluster definition files.

C.	You can use multiple compute resources on the same pipeline by defining multiple clusters for your activity.
D.	You cannot use multiple compute resources on the same pipeline.

Answer: A
Explanation:
Multiple compute resources can be used on the same pipeline in AWS Data Pipeline by defining multiple cluster objects in your definition file and associating the cluster to use for each activity via its runsOn field, which allows pipelines to combine AWS and on-premise resources, or to use a mix of instance types for their activities.
https://aws.amazon.com/datapipeline/faqs/

QUESTION 125
The two policies that you attach to an IAM role are the access policy and the trust policy. The trust policy identifies who can assume the role and grants the permission in the AWS Lambda account principal by adding the _____ action.

A.	aws:AssumeAdmin
B.	lambda:InvokeAsync
C.	sts:InvokeAsync
D.	sts:AssumeRole

Answer: D
Explanation:
The two policies that you attach to an IAM role are the access policy and the trust policy. Remember that adding an account to the trust policy of a role is only half of establishing the trust relationship. By default, no users in the trusted accounts can assume the role until the administrator for that account grants the users the permission to assume the role by adding the Amazon Resource Name (ARN) of the role to an Allow element for the sts:AssumeRole action.
http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_manage_modify.html

QUESTION 126
The MySecureData company has five branches across the globe. They want to expand their data centers such that their web server will be in the AWS and each branch would have their own database in the local data center. Based on the user login, the company wants to connect to the data center.
How can MySecureData company implement this scenario with the AWS VPC?

A.	Create five VPCs with the public subnet for the app server and setup the VPN gateway for each VPN to connect them individually.
B.	Use the AWS VPN CloudHub to communicate with multiple VPN connections.
C.	Use the AWS CloudGateway to communicate with multiple VPN connections.
D.	It is not possible to connect different data centers from a single VPC.

Answer: B
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. If the organization has multiple VPN connections, he can provide secure communication between sites using the AWS VPN CloudHub.
The VPN CloudHub operates on a simple hub-and-spoke model that the user can use with or

without a VPC. This design is suitable for customers with multiple branch offices and existing internet connections who would like to implement a convenient, potentially low-cost hub-and-spoke model for primary or backup connectivity between remote offices.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPN_CloudHub.html

QUESTION 127
One of your AWS Data Pipeline activities has failed consequently and has entered a hard failure state after retrying thrice.
You want to try it again. Is it possible to increase the number of automatic retries to more than thrice?

A.	Yes, you can increase the number of automatic retries to 6.
B.	Yes, you can increase the number of automatic retries to indefinite number.
C.	No, you cannot increase the number of automatic retries.
D.	Yes, you can increase the number of automatic retries to 10.

Answer: D
Explanation:
In AWS Data Pipeline, an activity fails if all of its activity attempts return with a failed state. By default, an activity retries three times before entering a hard failure state. You can increase the number of automatic retries to 10. However, the system does not allow indefinite retries.
https://aws.amazon.com/datapipeline/faqs/

QUESTION 128
True or False: In Amazon ElastiCache replication groups of Redis, for performance tuning reasons, you can change the roles of the cache nodes within the replication group, with the primary and one of the replicas exchanging roles.

A.	True, however, you get lower performance.
B.	FALSE
C.	TRUE
D.	False, you must recreate the replication group to improve performance tuning.

Answer: C
Explanation:
In Amazon ElastiCache, a replication group is a collection of Redis Cache Clusters, with one primary read-write cluster and up to five secondary, read-only clusters, which are called read replicas. You can change the roles of the cache clusters within the replication group, with the primary cluster and one of the replicas exchanging roles. You might decide to do this for performance tuning reasons.
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/Replication.Redis.Groups.html

QUESTION 129
How much memory does the cr1.8xlarge instance type provide?

A.	224 GB
B.	124 GB
C.	184 GB
D.	244 GB

Answer: D
Explanation:
The CR1 instances are part of the memory optimized instances. They offer lowest cost per GB RAM among all the AWS instance families. CR1 instances are part of the new generation of memory optimized instances, which can offer up to 244 GB RAM and run on faster CPUs (Intel Xeon E5-2670 with NUMA support) in comparison to the M2 instances of the same family. They support cluster networking for bandwidth intensive applications. cr1.8xlarge is one of the largest instance types of the CR1 family, which can offer 244 GB RAM.
http://aws.amazon.com/ec2/instance-types/

QUESTION 130
How many cg1.4xlarge on-demand instances can a user run in one region without taking any limit increase approval from AWS?

A.	20
B.	2
C.	5
D.	10

Answer: B
Explanation:
Generally AWS EC2 allows running 20 on-demand instances and 100 spot instances at a time. This limit can be increased by requesting at https://aws.amazon.com/contact-us/ec2-request. Excluding certain types of instances, the limit is lower than mentioned above. For cg1.4xlarge, the user can run only 2 on-demand instances at a time.
http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_ec2

QUESTION 131
Regarding Amazon SNS, you can send notification messages to mobile devices through any of the following supported push notification services, EXCEPT:

A.	Microsoft Windows Mobile Messaging (MWMM)
B.	Google Cloud Messaging for Android (GCM)
C.	Amazon Device Messaging (ADM)
D.	Apple Push Notification Service (APNS)

Answer: A
Explanation:
In Amazon SNS, you have the ability to send notification messages directly to apps on mobile devices. Notification messages sent to a mobile endpoint can appear in the mobile app as message alerts, badge updates, or even sound alerts. Microsoft Windows Mobile Messaging (MWMM) doesn't exist and is not supported by Amazon SNS.
http://docs.aws.amazon.com/sns/latest/dg/SNSMobilePush.html

QUESTION 132
You want to define permissions for a role in an IAM policy. Which of the following configuration formats should you use?

A.	An XML document written in the IAM Policy Language
B.	An XML document written in a language of your choice
C.	A JSON document written in the IAM Policy Language

D.	A JSON document written in a language of your choice

Answer: C
Explanation:
You define the permissions for a role in an IAM policy. An IAM policy is a JSON document written in the IAM Policy Language.
http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_terms-and-concepts.html

QUESTION 133
IAM Secure And Scalable is an organization which provides scalable and secure SAAS to its clients. They are planning to host a web server and App server on AWS VPC as separate tiers. The organization wants to implement the scalability by configuring Auto Scaling and load balancer with their app servers (middle tier) too.
Which of the below mentioned options suits their requirements?

A.	Since ELB is internet facing, it is recommended to setup HAProxy as the Load balancer within the VPC.
B.	Create an Internet facing ELB with VPC and configure all the App servers with it.
C.	The user should make ELB with EC2-CLASSIC and enable SSH with it for security.
D.	Create an Internal Load balancer with VPC and register all the App servers with it.

Answer: D
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances.
There are two ELBs available with VPC: internet facing and internal (private) ELB. For internal servers, such as App servers the organization can create an internal load balancer in their VPC and then place back-end application instances behind the internal load balancer. The internal load balancer will route requests to the back-end application instances, which are also using private IP addresses and only accept requests from the internal load balancer.
http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/vpc-loadbalancer-types.html

QUESTION 134
True or False: Amazon ElastiCache supports the Redis key-value store.

A.	True, ElastiCache supports the Redis key-value store, but with limited functionalities.
B.	False, ElastiCache does not support the Redis key-value store.
C.	True, ElastiCache supports the Redis key-value store.
D.	False, ElastiCache supports the Redis key-value store only if you are in a VPC environment.

Answer: C
Explanation:
This is true. ElastiCache supports two open-source in-memory caching engines: 1. Memcached - a widely adopted memory object caching system. ElastiCache is protocol compliant with Memcached, so popular tools that you use today with existing Memcached environments will work seamlessly with the service. 2. Redis - a popular open-source in-memory key-value store that supports data structures such as sorted sets and lists. ElastiCache supports Master / Slave replication and Multi-AZ which can be used to achieve cross AZ redundancy.
https://aws.amazon.com/elasticache/

QUESTION 135
Which of the following is NOT an advantage of using AWS Direct Connect?

A.	AWS Direct Connect provides users access to public and private resources by using two different connections while maintaining network separation between the public and private environments.
B.	AWS Direct Connect provides a more consistent network experience than Internet-based connections.
C.	AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS.
D.	AWS Direct Connect reduces your network costs.

Answer: A
Explanation:
AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your datacenter, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections.
By using industry standard 802.1q VLANs, this dedicated connection can be partitioned into multiple virtual interfaces. This allows you to use the same connection to access public resources such as objects stored in Amazon S3 using public IP address space, and private resources such as Amazon EC2 instances running within an Amazon Virtual Private Cloud (VPC) using private IP space, while maintaining network separation between the public and private environments.
http://aws.amazon.com/directconnect/#details

QUESTION 136
An organization is setting up an application on AWS to have both High Availabilty (HA) and Disaster Recovery (DR). The organization wants to have both Recovery point objective (RPO) and Recovery time objective (RTO) of 10 minutes. Which of the below mentioned service configurations does not help the organization achieve the said RPO and RTO?

A.	Take a snapshot of the data every 10 minutes and copy it to the other region.
B.	Use an elastic IP to assign to a running instance and use Route 53 to map the user's domain with that IP.
C.	Create ELB with multi-region routing to allow automated failover when required.
D.	Use an AMI copy to keep the AMI available in other regions.

Answer: C
Explanation:
AWS provides an on demand, scalable infrastructure. AWS EC2 allows the user to launch On-Demand instances and the organization should create an AMI of the running instance. Copy the AMI to another region to enable Disaster Recovery (DR) in case of region failure. The organization should also use EBS for persistent storage and take a snapshot every 10 minutes to meet Recovery time objective (RTO). They should also setup an elastic IP and use it with Route 53 to route requests to the same IP. When one of the instances fails the organization can launch new instances and assign the same EIP to a new instance to achieve High Availability (HA). The ELB works only for a particular region and does not route requests across regions.
http://d36cz9buwru1tt.cloudfront.net/AWS_Disaster_Recovery.pdf

QUESTION 137
An organization is having an application which can start and stop an EC2 instance as per

schedule. The organization needs the MAC address of the instance to be registered with its software. The instance is launched in EC2-CLASSIC. How can the organization update the MAC registration every time an instance is booted?

A.	The organization should write a boot strapping script which will get the MAC address from the instance metadata and use that script to register with the application.
B.	The organization should provide a MAC address as a part of the user data. Thus, whenever the instance is booted the script assigns the fixed MAC address to that instance.
C.	The instance MAC address never changes. Thus, it is not required to register the MAC address every time.
D.	AWS never provides a MAC address to an instance; instead the instance ID is used for identifying the instance for any software registration.

Answer: A
Explanation:
AWS provides an on demand, scalable infrastructure. AWS EC2 allows the user to launch On-Demand instances. AWS does not provide a fixed MAC address to the instances launched in EC2-CLASSIC. If the instance is launched as a part of EC2-VPC, it can have an ENI which can have a fixed MAC. However, with EC2-CLASSIC, every time the instance is started or stopped it will have a new MAC address. To get this MAC, the organization can run a script on boot which can fetch the instance metadata and get the MAC address from that instance metadata. Once the MAC is received, the organization can register that MAC with the software.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AESDG-chapter-instancedata.html

QUESTION 138
Does Amazon RDS API provide actions to modify DB instances inside a VPC and associate them with DB Security Groups?

A.	Yes, Amazon does this but only for MySQL RDS.
B.	Yes
C.	No
D.	Yes, Amazon does this but only for Oracle RDS.

Answer: B
Explanation:
You can use the action Modify DB Instance, available in the Amazon RDS API, to pass values for the parameters DB Instance Identifier and DB Security Groups specifying the instance ID and the DB Security Groups you want your instance to be part of.
http://docs.aws.amazon.com/AmazonRDS/latest/APIReference/API_ModifyDBInstance.html

QUESTION 139
An organization is setting up a backup and restore system in AWS of their in premise system. The organization needs High Availability(HA) and Disaster Recovery(DR) but is okay to have a longer recovery time to save costs. Which of the below mentioned setup options helps achieve the objective of cost saving as well as DR in the most effective way?

A.	Setup pre- configured servers and create AMIs.. Use EIP and Route 53 to quickly switch over to AWS from in premise.
B.	Setup the backup data on S3 and transfer data to S3 regularly using the storage gateway.
C.	Setup a small instance with AutoScaling; in case of DR start diverting all the load to AWS from on premise.
D.	Replicate on premise DB to EC2 at regular intervals and setup a scenario similar to the pilot light.

Answer: B
Explanation:
AWS has many solutions for Disaster Recovery(DR) and High Availability(HA). When the organization wants to have HA and DR but are okay to have a longer recovery time they should select the option backup and restore with S3. The data can be sent to S3 using either Direct Connect, Storage Gateway or over the internet.
The EC2 instance will pick the data from the S3 bucket when started and setup the environment. This process takes longer but is very cost effective due to the low pricing of S3. In all the other options, the EC2 instance might be running or there will be AMI storage costs. Thus, it will be a costlier option. In this scenario the organization should plan appropriate tools to take a backup, plan the retention policy for data and setup security of the data.
http://d36cz9buwru1tt.cloudfront.net/AWS_Disaster_Recovery.pdf

QUESTION 140
By default, what is the maximum number of Cache Nodes you can run in Amazon ElastiCache?

A.	20
B.	50
C.	100
D.	200

Answer: A
Explanation:
In Amazon ElastiCache, you can run a maximum of 20 Cache Nodes.
http://aws.amazon.com/elasticache/faqs/

QUESTION 141
Does an AWS Direct Connect location provide access to Amazon Web Services in the region it is associated with as well as access to other US regions?

A.	No, it provides access only to the region it is associated with.
B.	No, it provides access only to the US regions other than the region it is associated with.
C.	Yes, it provides access.
D.	Yes, it provides access but only when there's just one Availability Zone in the region.

Answer: C
Explanation:
An AWS Direct Connect location provides access to Amazon Web Services in the region it is associated with, as well as access to other US regions. For example, you can provision a single connection to any AWS Direct Connect location in the US and use it to access public AWS services in all US Regions and AWS GovCloud (US).
http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

QUESTION 142
Which of the following components of AWS Data Pipeline specifies the business logic of your data management?

A.	Task Runner
B.	Pipeline definition
C.	AWS Direct Connect

D.	Amazon Simple Storage Service (Amazon S3)

Answer: B
Explanation:
A pipeline definition specifies the business logic of your data management.
http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/what-is-datapipeline.html

QUESTION 143
What feature of the load balancing service attempts to force subsequent connections to a service to be redirected to the same node as long as it is online?

A.	Node balance
B.	Session retention
C.	Session multiplexing
D.	Session persistence

Answer: D
Explanation:
Session persistence is a feature of the load balancing service. It attempts to force subsequent connections to a service to be redirected to the same node as long as it is online.
http://docs.rackspace.com/loadbalancers/api/v1.0/clb-devguide/content/Concepts-d1e233.html

QUESTION 144
What types of identities do Amazon Cognito identity pools support?

A.	They support both authenticated and unauthenticated identities.
B.	They support only unauthenticated identities.
C.	They support neither authenticated nor unauthenticated identities.
D.	They support only authenticated identities.

Answer: A
Explanation:
Amazon Cognito identity pools support both authenticated and unauthenticated identities. Authenticated identities belong to users who are authenticated by a public login provider or your own backend authentication process. Unauthenticated identities typically belong to guest users.
http://docs.aws.amazon.com/cognito/devguide/identity/identity-pools/

QUESTION 145
In IAM, which of the following is true of temporary security credentials?

A.	Once you issue temporary security credentials, they cannot be revoked.
B.	None of these are correct.
C.	Once you issue temporary security credentials, they can be revoked only when the virtual MFA device is used.
D.	Once you issue temporary security credentials, they can be revoked.

Answer: A
Explanation:
Temporary credentials in IAM are valid throughout their defined duration of time and hence can't be revoked. However, because permissions are evaluated each time an AWS request is made using the credentials, you can achieve the effect of revoking the credentials by changing the

permissions for the credentials even after they have been issued.
http://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp_control-access_disable-perms.h tml

QUESTION 146
The CFO of a company wants to allow one of his employees to view only the AWS usage report page.
Which of the below mentioned IAM policy statements allows the user to have access to the AWS usage report page?

A.	"Effect": "Allow", "Action": ["Describe"], "Resource": "Billing"
B.	"Effect": "Allow", "Action": ["aws-portal: ViewBilling"], "Resource": "*"
C.	"Effect": "Allow", "Action": ["aws-portal:ViewUsage"], "Resource": "*"
D.	"Effect": "Allow", "Action": ["AccountUsage], "Resource": "*"

Answer: C
Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the CFO wants to allow only AWS usage report page access, the policy for that IAM user will be as given below:
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
"Action": [
"aws-portal:ViewUsage"
],
"Resource": "*"
}
]
}
http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/billing-permissions-ref.html

QUESTION 147
In Amazon VPC, what is the default maximum number of BGP advertised routes allowed per route table?

A.	15
B.	100
C.	5
D.	10

Answer: B
Explanation:
The maximum number of BGP advertised routes allowed per route table is 100.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Appendix_Limits.html

QUESTION 148
An organization has created 5 IAM users. The organization wants to give them the same login ID but different passwords. How can the organization achieve this?

A.	The organization should create each user in a separate region so that they have their own URL to login
B.	The organization should create a separate login ID but give the IAM users the same alias so that each one can login with their alias
C.	It is not possible to have the same login ID for multiple IAM users of the same account
D.	The organization should create various groups and add each user with the same login ID to different groups. The user can login with their own group ID

Answer: C
Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services.
Whenever the organization is creating an IAM user, there should be a unique ID for each user. It is not possible to have the same login ID for multiple users. The names of users, groups, roles, instance profiles must be alphanumeric, including the following common characters: plus (+), equal (=), comma (,), period (.), at (@), and dash (-).
http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_SettingUpUser.html

QUESTION 149
The user has provisioned the PIOPS volume with an EBS optimized instance.
Generally speaking, in which I/O chunk should the bandwidth experienced by the user be measured by AWS?

A.	128 KB
B.	256 KB
C.	64 KB
D.	32 KB

Answer: B
Explanation:
IOPS are input/output operations per second. Amazon EBS measures each I/O operation per second (that is 256 KB or smaller) as one IOPS.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-io-characteristics.html

QUESTION 150
A user is planning to use EBS for his DB requirement. The user already has an EC2 instance running in the VPC private subnet.
How can the user attach the EBS volume to a running instance?

A.	The user can create EBS in the same zone as the subnet of instance and attach that EBS to instance.
B.	It is not possible to attach an EBS to an instance running in VPC until the instance is stopped.
C.	The user can specify the same subnet while creating EBS and then attach it to a running instance.
D.	The user must create EBS within the same VPC and then attach it to a running instance.

Answer: A
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. The VPC is always specific to a region. The user can create a VPC which can span multiple Availability Zones by adding one or more subnets in each Availability Zone.

The instance launched will always be in the same availability zone of the respective subnet. When creating an EBS the user cannot specify the subnet or VPC. However, the user must create the EBS in the same zone as the instance so that it can attach the EBS volume to the running instance.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html#VPCSubnet

QUESTION 151
An organization is planning to setup a management network on the AWS VPC. The organization is trying to secure the webserver on a single VPC instance such that it allows the internet traffic as well as the back-end management traffic. The organization wants to make so that the back end management network interface can receive the SSH traffic only from a selected IP range, while the internet facing webserver will have an IP address which can receive traffic from all the internet IPs. How can the organization achieve this by running web server on a single instance?

A.	It is not possible to have two IP addresses for a single instance.
B.	The organization should create two network interfaces with the same subnet and security group to assign separate IPs to each network interface.
C.	The organization should create two network interfaces with separate subnets so one instance can have two subnets and the respective security groups for controlled access.
D.	The organization should launch an instance with two separate subnets using the same network interface which allows to have a separate CIDR as well as security groups.

Answer: C
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. An Elastic Network Interface (ENI) is a virtual network interface that the user can attach to an instance in a VPC. The user can create a management network using two separate network interfaces. For the present scenario it is required that the secondary network interface on the instance handles the public facing traffic and the primary network interface handles the back-end management traffic and it is connected to a separate subnet in the VPC that has more restrictive access controls. The public facing interface, which may or may not be behind a load balancer, has an associated security group to allow access to the server from the internet while the private facing interface has an associated security group allowing SSH access only from an allowed range of IP addresses either within the VPC or from the internet, a private subnet within the VPC or a virtual private gateway.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 152
A user is trying to create a vault in AWS Glacier. The user wants to enable notifications.
In which of the below mentioned options can the user enable the notifications from the AWS console?

A.	Glacier does not support the AWS console
B.	Archival Upload Complete
C.	Vault Upload Job Complete
D.	Vault Inventory Retrieval Job Complete

Answer: D
Explanation:
From AWS console the user can configure to have notifications sent to Amazon Simple Notifications Service (SNS). The user can select specific jobs that, on completion, will trigger the notifications such as Vault Inventory Retrieval Job Complete and Archive Retrieval Job Complete.

http://docs.aws.amazon.com/amazonglacier/latest/dev/configuring-notifications-console.html

QUESTION 153
An organization is purchasing licensed software. The software license can be registered only to a specific MAC Address. The organization is going to host the software in the AWS environment. How can the organization fulfil the license requirement as the MAC address changes every time an instance is started/stopped/terminated?

A.	It is not possible to have a fixed MAC address with AWS.
B.	The organization should use VPC with the private subnet and configure the MAC address with that subnet.
C.	The organization should use VPC with an elastic network interface which will have a fixed MAC Address.
D.	The organization should use VPC since VPC allows to configure the MAC address for each EC2 instance.

Answer: C
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. An Elastic Network Interface (ENI) is a virtual network interface that the user can attach to an instance in a VPC. An ENI can include attributes such as: a primary private IP address, one or more secondary private IP addresses, one elastic IP address per private IP address, one public IP address, one or more security groups, a MAC address, a source/destination check flag, and a description. The user can create a network interface, attach it to an instance, detach it from an instance, and attach it to another instance. The attributes of a network interface follow the network interface as it is attached or detached from an instance and reattached to another instance. Thus, the user can maintain a fixed MAC using the network interface.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 154
ExamKiller has three separate departments and each department has their own AWS accounts. The HR department has created a file sharing site where all the on roll employees' data is uploaded. The Admin department uploads data about the employee presence in the office to their DB hosted in the VPC. The Finance department needs to access data from the HR department to know the on roll employees to calculate the salary based on the number of days that an employee is present in the office.
How can ExamKiller setup this scenario?

A.	It is not possible to configure VPC peering since each department has a separate AWS account.
B.	Setup VPC peering for the VPCs of Admin and Finance.
C.	Setup VPC peering for the VPCs of Finance and HR as well as between the VPCs of Finance and Admin.
D.	Setup VPC peering for the VPCs of Admin and HR

Answer: C
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. A VPC peering connection allows the user to route traffic between the peer VPCs using private IP addresses as if they are a part of the same network. This is helpful when one VPC from the same or different AWS account wants to connect with resources of the other VPC.
http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/peering-configurations-full-

access.html#t hree-vpcs-full-access

QUESTION 155
An organization is undergoing a security audit. The auditor wants to view the AWS VPC configurations as the organization has hosted all the applications in the AWS VPC. The auditor is from a remote place and wants to have access to AWS to view all the VPC records.
How can the organization meet the expectations of the auditor without compromising on the security of their AWS infrastructure?

A.	The organization should not accept the request as sharing the credentials means compromising on security.
B.	Create an IAM role which will have read only access to all EC2 services including VPC and assign that role to the auditor .
C.	Create an IAM user who will have read only access to the AWS VPC and share those credentials with the auditor.
D.	The organization should create an IAM user with VPC full access but set a condition that will not allow to modify anything if the request is from any IP other than the organization's data center.

Answer: C
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. The VPC also works with IAM and the organization can create IAM users who have access to various VPC services. If an auditor wants to have access to the AWS VPC to verify the rules, the organization should be careful before sharing any data which can allow making updates to the AWS infrastructure. In this scenario it is recommended that the organization creates an IAM user who will have read only access to the VPC. Share the above mentioned credentials with the auditor as it cannot harm the organization.
The sample policy is given below:
{
"Effect":"Allow",
"Action":[
"ec2:DescribeVpcs",
"ec2:DescribeSubnets",
"ec2:DescribeInternetGateways",
"ec2:DescribeCustomerGateways",
"ec2:DescribeVpnGateways",
"ec2:DescribeVpnConnections",
"ec2:DescribeRouteTables",
"ec2:DescribeAddresses",
"ec2:DescribeSecurityGroups",
"ec2:DescribeNetworkAcls",
"ec2:DescribeDhcpOptions",
"ec2:DescribeTags",
"ec2:DescribeInstances"
],
"Resource":"*"
}
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_IAM.html

QUESTION 156
What is the maximum length for an instance profile name in AWS IAM?

A.	512 characters

B.	128 characters
C.	1024 characters
D.	64 characters

Answer: B
Explanation:
The maximum length for an instance profile name is 128 characters.
http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html

QUESTION 157
Cognito Sync is an AWS service that you can use to synchronize user profile data across mobile devices without requiring your own backend. When the device is online, you can synchronize data. If you also set up push sync, what does it allow you to do?

A.	Notify other devices that a user profile is available across multiple devices
B.	Synchronize user profile data with less latency
C.	Notify other devices immediately that an update is available
D.	Synchronize online data faster

Answer: C
Explanation:
Cognito Sync is an AWS service that you can use to synchronize user profile data across mobile devices without requiring your own backend. When the device is online, you can synchronize data, and if you have also set up push sync, notify other devices immediately that an update is available.
http://docs.aws.amazon.com/cognito/devguide/sync/

QUESTION 158
An organization is planning to create a secure scalable application with AWS VPC and ELB. The organization has two instances already running and each instance has an ENI attached to it in addition to a primary network interface. The primary network interface and additional ENI both have an elastic IP attached to it.
If those instances are registered with ELB and the organization wants ELB to send data to a particular EIP of the instance, how can they achieve this?

A.	The organization should ensure that the IP which is required to receive the ELB traffic is attached to a primary network interface.
B.	It is not possible to attach an instance with two ENIs with ELB as it will give an IP conflict error.
C.	The organization should ensure that the IP which is required to receive the ELB traffic is attached to an additional ENI.
D.	It is not possible to send data to a particular IP as ELB will send to any one EIP.

Answer: A
Explanation:
Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances. There are two ELBs available with VPC: internet facing and internal (private) ELB. For the internet facing ELB it is required that the ELB should be in a public subnet. When the user registers a multi-homed instance (an instance that has an Elastic Network Interface (ENI) attached) with a load balancer, the load balancer will route the traffic to the IP address of the primary network interface (eth0).

http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/gs-ec2VPC.html

QUESTION 159
In Amazon Cognito, your mobile app authenticates with the Identity Provider (IdP) using the provider's SDK. Once the end user is authenticated with the IdP, the OAuth or OpenID Connect token returned from the IdP is passed by your app to Amazon Cognito, which returns a new _____ for the user and a set of temporary, limited-privilege AWS credentials.

A.	Cognito Key Pair
B.	Cognito API
C.	Cognito ID
D.	Cognito SDK

Answer: C
Explanation:
Your mobile app authenticates with the identity provider (IdP) using the provider's SDK. Once the end user is authenticated with the IdP, the OAuth or OpenID Connect token returned from the IdP is passed by your app to Amazon Cognito, which returns a new Cognito ID for the user and a set of temporary, limited-privilege AWS credentials.
http://aws.amazon.com/cognito/faqs/

QUESTION 160
What is the maximum length for a certificate ID in AWS IAM?

A.	1024 characters
B.	512 characters
C.	64 characters
D.	128 characters

Answer: D
Explanation:
The maximum length for a certificate ID is 128 characters.
http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html

QUESTION 161
A user is trying to create a PIOPS EBS volume with 3 GB size and 90 IOPS. Will AWS create the volume?

A.	No, since the PIOPS and EBS size ratio is less than 30
B.	Yes, since the ratio between EBS and IOPS is less than 30
C.	No, the EBS size is less than 4GB
D.	Yes, since PIOPS is higher than 100

Answer: C
Explanation:
A Provisioned IOPS (SSD) volume can range in size from 4 GiB to 16 TiB and you can provision up to 20,000 IOPS per volume.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html#EBSVolumeTypes_piops

QUESTION 162
A user has configured EBS volume with PIOPS. The user is not experiencing the optimal throughput. Which of the following could not be factor affecting I/O performance of that EBS volume?

A.	EBS bandwidth of dedicated instance exceeding the PIOPS
B.	EC2 bandwidth
C.	EBS volume size
D.	Instance type is not EBS optimized

Answer: C
Explanation:
If the user is not experiencing the expected IOPS or throughput that is provisioned, ensure that the EC2 bandwidth is not the limiting factor, the instance is EBS-optimized (or include 10 Gigabit network connectivity) and the instance type EBS dedicated bandwidth exceeds the IOPS more than he has provisioned.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-io-characteristics.html

QUESTION 163
If a single condition within an IAM policy includes multiple values for one key, it will be evaluated using a logical _____.

A.	OR
B.	NAND
C.	NOR
D.	AND

Answer: A
Explanation:
If a single condition within an IAM policy includes multiple values for one key, it will be evaluated using a logical OR.
http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html

QUESTION 164
Which of the following cache engines does Amazon ElastiCache support?

A.	Amazon ElastiCache supports Memcached and Redis.
B.	Amazon ElastiCache supports Redis and WinCache.
C.	Amazon ElastiCache supports Memcached and Hazelcast.
D.	Amazon ElastiCache supports Memcached only.

Answer: A
Explanation:
The cache engines supported by Amazon ElastiCache are Memcached and Redis.
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/SelectEngine.html

QUESTION 165
You have been given the task to define multiple AWS Data Pipeline schedules for different activities in the same pipeline. Which of the following would successfully accomplish this task?

A.	Creating multiple pipeline definition files

B.	Defining multiple pipeline definitions in your schedule objects file and associating the desired schedule to the correct activity via its schedule field
C.	Defining multiple schedule objects in your pipeline definition file and associating the desired schedule to the correct activity via its schedule field
D.	Defining multiple schedule objects in the schedule field

Answer: C
Explanation:
To define multiple schedules for different activities in the same pipeline, in AWS Data Pipeline, you should define multiple schedule objects in your pipeline definition file and associate the desired schedule to the correct activity via its schedule field. As an example of this, it could allow you to define a pipeline in which log files are stored in Amazon S3 each hour to drive generation of an aggregate report once a day.
https://aws.amazon.com/datapipeline/faqs/

QUESTION 166
In a VPC, can you modify a set of DHCP options after you create them?

A.	Yes, you can modify a set of DHCP options within 48 hours after creation and there are no VPCs associated with them.
B.	Yes, you can modify a set of DHCP options any time after you create them.
C.	No, you can't modify a set of DHCP options after you create them.
D.	Yes, you can modify a set of DHCP options within 24 hours after creation.

Answer: C
Explanation:
After you create a set of DHCP options, you can't modify them. If you want your VPC to use a different set of DHCP options, you must create a new set and associate them with your VPC. You can also set up your VPC to use no DHCP options at all.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_DHCP_Options.html

QUESTION 167
A bucket owner has allowed another account's IAM users to upload or access objects in his bucket. The IAM user of Account A is trying to access an object created by the IAM user of account B. What will happen in this scenario?

A.	It is not possible to give permission to multiple IAM users
B.	AWS S3 will verify proper rights given by the owner of Account A, the bucket owner as well as by the IAM user B to the object
C.	The bucket policy may not be created as S3 will give error due to conflict of Access Rights
D.	It is not possible that the IAM user of one account accesses objects of the other IAM user

Answer: B
Explanation:
If a IAM user is trying to perform some action on an object belonging to another AWS user's bucket, S3 will verify whether the owner of the IAM user has given sufficient permission to him. It also verifies the policy for the bucket as well as the policy defined by the object owner.
http://docs.aws.amazon.com/AmazonS3/latest/dev/access-control-auth-workflow-object-operation.html

QUESTION 168

Which statement is NOT true about a stack which has been created in a Virtual Private Cloud (VPC) in AWS OpsWorks?

A.	Subnets whose instances cannot communicate with the Internet are referred to as public subnets.
B.	Subnets whose instances can communicate only with other instances in the VPC and cannot communicate directly with the Internet are referred to as private subnets.
C.	All instances in the stack should have access to any package repositories that your operating system depends on, such as the Amazon Linux or Ubuntu Linux repositories.
D.	Your app and custom cookbook repositories should be accessible for all instances in the stack.

Answer: A
Explanation:
In AWS OpsWorks, you can control user access to a stack's instances by creating it in a virtual private cloud (VPC). For example, you might not want users to have direct access to your stack's app servers or databases and instead require that all public traffic be channeled through an Elastic Load Balancer. A VPC consists of one or more subnets, each of which contains one or more instances. Each subnet has an associated routing table that directs outbound traffic based on its destination IP address. Instances within a VPC can generally communicate with each other, regardless of their subnet. Subnets whose instances can communicate with the Internet are referred to as public subnets. Subnets whose instances can communicate only with other instances in the VPC and cannot communicate directly with the Internet are referred to as private subnets. AWS OpsWorks requires the VPC to be configured so that every instance in the stack, including instances in private subnets, has access to the following endpoints:
The AWS OpsWorks service, https://opsworks-instance-service.us-east-1.amazonaws.com .
Amazon S3
The package repositories for Amazon Linux or Ubuntu 12.04 LTS, depending on which operating system you specify.
Your app and custom cookbook repositories.
http://docs.aws.amazon.com/opsworks/latest/userguide/workingstacks-vpc.html#workingstacks-vpc-basi cs

QUESTION 169
An organization has hosted an application on the EC2 instances. There will be multiple users connecting to the instance for setup and configuration of application. The organization is planning to implement certain security best practices. Which of the below mentioned pointers will not help the organization achieve better security arrangement?

A.	Allow only IAM users to connect with the EC2 instances with their own secret access key.
B.	Create a procedure to revoke the access rights of the individual user when they are not required to connect to EC2 instance anymore for the purpose of application configuration.
C.	Apply the latest patch of OS and always keep it updated.
D.	Disable the password based login for all the users. All the users should use their own keys to connect with the instance securely.

Answer: A
Explanation:
Since AWS is a public cloud any application hosted on EC2 is prone to hacker attacks. It becomes extremely important for a user to setup a proper security mechanism on the EC2 instances. A few of the security measures are listed below:
Always keep the OS updated with the latest patch
Always create separate users with in OS if they need to connect with the EC2 instances, create their keys and disable their password
Create a procedure using which the admin can revoke the access of the user when the business work on the EC2 instance is completed

Lock down unnecessary ports
Audit any proprietary applications that the user may be running on the EC2 instance Provide temporary escalated privileges, such as sudo for users who need to perform occasional privileged tasks
The IAM is useful when users are required to work with AWS resources and actions, such as launching an instance. It is not useful to connect (RDP / SSH) with an instance.
http://aws.amazon.com/articles/1233/

QUESTION 170
By default, temporary security credentials for an IAM user are valid for a maximum of 12 hours, but you can request a duration as long as _____ hours.

A.	24
B.	36
C.	10
D.	48

Answer: B
Explanation:
By default, temporary security credentials for an IAM user are valid for a maximum of 12 hours, but you can request a duration as short as 15 minutes or as long as 36 hours.
http://docs.aws.amazon.com/STS/latest/UsingSTS/CreatingSessionTokens.html

QUESTION 171
What RAID method is used on the Cloud Block Storage back-end to implement a very high level of reliability and performance?

A.	RAID 1 (Mirror)
B.	RAID 5 (Blocks striped, distributed parity)
C.	RAID 10 (Blocks mirrored and striped)
D.	RAID 2 (Bit level striping)

Answer: C
Explanation:
Cloud Block Storage back-end storage volumes employs the RAID 10 method to provide a very high level of reliability and performance.
http://www.rackspace.com/knowledge_center/product-faq/cloud-block-storage

QUESTION 172
One of the AWS account owners faced a major challenge in June as his account was hacked and the hacker deleted all the data from his AWS account. This resulted in a major blow to the business. Which of the below mentioned steps would not have helped in preventing this action?

A.	Setup an MFA for each user as well as for the root account user.
B.	Take a backup of the critical data to offsite / on premise.
C.	Create an AMI and a snapshot of the data at regular intervals as well as keep a copy to separate regions.
D.	Do not share the AWS access and secret access keys with others as well do not store it inside programs, instead use IAM roles.

Answer: C

Explanation:
AWS security follows the shared security model where the user is as much responsible as Amazon. If the user wants to have secure access to AWS while hosting applications on EC2, the first security rule to follow is to enable MFA for all users. This will add an added security layer. In the second step, the user should never give his access or secret access keys to anyone as well as store inside programs. The better solution is to use IAM roles. For critical data of the organization, the user should keep an offsite/ in premise backup which will help to recover critical data in case of security breach. It is recommended to have AWS AMIs and snapshots as well as keep them at other regions so that they will help in the DR scenario. However, in case of a data security breach of the account they may not be very helpful as hacker can delete that.
Therefore ,creating an AMI and a snapshot of the data at regular intervals as well as keep a copy to separate regions, would not have helped in preventing this action.
http://media.amazonwebservices.com/pdf/AWS_Security_Whitepaper.pdf

QUESTION 173
With Amazon Elastic MapReduce (Amazon EMR) you can analyze and process vast amounts of data. The cluster is managed using an open-source framework called Hadoop. You have set up an application to run Hadoop jobs. The application reads data from DynamoDB and generates a temporary file of 100 TBs.
The whole process runs for 30 minutes and the output of the job is stored to S3. Which of the below mentioned options is the most cost effective solution in this case?

A.	Use Spot Instances to run Hadoop jobs and configure them with EBS volumes for persistent data storage.
B.	Use Spot Instances to run Hadoop jobs and configure them with ephermal storage for output file storage.
C.	Use an on demand instance to run Hadoop jobs and configure them with EBS volumes for persistent storage.
D.	Use an on demand instance to run Hadoop jobs and configure them with ephemeral storage for output file storage.

Answer: B
Explanation:
AWS EC2 Spot Instances allow the user to quote his own price for the EC2 computing capacity. The user can simply bid on the spare Amazon EC2 instances and run them whenever his bid exceeds the current Spot Price. The Spot Instance pricing model complements the On-Demand and Reserved Instance pricing models, providing potentially the most cost-effective option for obtaining compute capacity, depending on the application. The only challenge with a Spot Instance is data persistence as the instance can be terminated whenever the spot price exceeds the bid price. In the current scenario a Hadoop job is a temporary job and does not run for a longer period. It fetches data from a persistent DynamoDB. Thus, even if the instance gets terminated there will be no data loss and the job can be re-run. As the output files are large temporary files, it will be useful to store data on ephermal storage for cost savings.
http://aws.amazon.com/ec2/purchasing-options/spot-instances/

QUESTION 174
In Amazon SNS, to send push notifications to mobile devices using Amazon SNS and ADM, you need to obtain the following, except:

A.	Device token
B.	Client ID
C.	Registration ID
D.	Client secret

Answer: A
Explanation:
To send push notifications to mobile devices using Amazon SNS and ADM, you need to obtain the following: Registration ID and Client secret.
http://docs.aws.amazon.com/sns/latest/dg/SNSMobilePushPrereq.html

QUESTION 175
True or False : "In the context of Amazon ElastiCache, from the application's point of view, connecting to the cluster configuration endpoint is no different than connecting directly to an individual cache node."

A.	True, from the application's point of view, connecting to the cluster configuration endpoint is no different than connecting directly to an individual cache node since, each has a unique node identifier.
B.	True, from the application's point of view, connecting to the cluster configuration endpoint is no different than connecting directly to an individual cache node.
C.	False, you can connect to a cache node, but not to a cluster configuration endpoint.
D.	False, you can connect to a cluster configuration endpoint, but not to a cache node.

Answer: B
Explanation:
This is true. From the application's point of view, connecting to the cluster configuration endpoint is no different than connecting directly to an individual cache node. In the process of connecting to cache nodes, the application resolves the configuration endpoint's DNS name. Because the configuration endpoint maintains CNAME entries for all of the cache nodes, the DNS name resolves to one of the nodes; the client can then connect to that node.
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/AutoDiscovery.HowAutoDiscoveryW orks.html

QUESTION 176
An organization is setting up a highly scalable application using Elastic Beanstalk.
They are using Elastic Load Balancing (ELB) as well as a Virtual Private Cloud (VPC) with public and private subnets. They have the following requirements:

```
- All the EC2 instances should have a private IP
- All the EC2 instances should receive data via the ELB's.
```

Which of these will not be needed in this setup?

A.	Launch the EC2 instances with only the public subnet.
B.	Create routing rules which will route all inbound traffic from ELB to the EC2 instances.
C.	Configure ELB and NAT as a part of the public subnet only.
D.	Create routing rules which will route all outbound traffic from the EC2 instances through NAT.

Answer: A
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. If the organization wants the Amazon EC2 instances to have a private IP address, he should create a public and private subnet for VPC in each Availability Zone (this is an AWS Elastic Beanstalk requirement). The organization should add their public resources, such as ELB and NAT to the public subnet, and

AWS Elastic Beanstalk will assign them unique elastic IP addresses (a static, public IP address). The organization should launch Amazon EC2 instances in a private subnet so that AWS Elastic Beanstalk assigns them non-routable private IP addresses. Now the organization should configure route tables with the following rules:
. route all inbound traffic from ELB to EC2 instances
. route all outbound traffic from EC2 instances through NAT
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/AWSHowTo-vpc.html

QUESTION 177
An EC2 instance that performs source/destination checks by default is launched in a private VPC subnet. All security, NACL, and routing definitions are configured as expected. A custom NAT instance is launched.
Which of the following must be done for the custom NAT instance to work?

A.	The source/destination checks should be disabled on the NAT instance.
B.	The NAT instance should be launched in public subnet.
C.	The NAT instance should be configured with a public IP address.
D.	The NAT instance should be configured with an elastic IP address.

Answer: A
Explanation:
Each EC2 instance performs source/destination checks by default. This means that the instance must be the source or destination of any traffic it sends or receives. However, a NAT instance must be able to send and receive traffic when the source or destination is not itself. Therefore, you must disable source/destination checks on the NAT instance.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_NAT_Instance.html#EIP_Disable_Src DestCheck

QUESTION 178
An organization has created multiple components of a single application for compartmentalization. Currently all the components are hosted on a single EC2 instance. Due to security reasons the organization wants to implement two separate SSLs for the separate modules although it is already using VPC. How can the organization achieve this with a single instance?

A.	You have to launch two instances each in a separate subnet and allow VPC peering for a single IP.
B.	Create a VPC instance which will have multiple network interfaces with multiple elastic IP addresses.
C.	Create a VPC instance which will have both the ACL and the security group attached to it and have separate rules for each IP address.
D.	Create a VPC instance which will have multiple subnets attached to it and each will have a separate IP address.

Answer: B
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. With VPC the user can specify multiple private IP addresses for his instances.
The number of network interfaces and private IP addresses that a user can specify for an instance depends on the instance type. With each network interface the organization can assign an EIP. This scenario helps when the user wants to host multiple websites on a single EC2 instance by using multiple SSL certificates on a single server and associating each certificate with

a specific EIP address. It also helps in scenarios for operating network appliances, such as firewalls or load balancers that have multiple private IP addresses for each network interface.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/MultipleIP.html

QUESTION 179
An organization is making software for the CIA in USA. CIA agreed to host the application on AWS but in a secure environment. The organization is thinking of hosting the application on the AWS GovCloud region. Which of the below mentioned difference is not correct when the organization is hosting on the AWS GovCloud in comparison with the AWS standard region?

A.	The billing for the AWS GovCLoud will be in a different account than the Standard AWS account.
B.	GovCloud region authentication is isolated from Amazon.com.
C.	Physical and logical administrative access only to U.S. persons.
D.	It is physically isolated and has logical network isolation from all the other regions.

Answer: A
Explanation:
AWS GovCloud (US) is an isolated AWS region designed to allow U.S. government agencies and customers to move sensitive workloads into the cloud by addressing their specific regulatory and compliance requirements. The AWS GovCloud (US) Region adheres to the U.S. International Traffic in Arms Regulations (ITAR) requirements. It has added advantages, such as:
Restricting physical and logical administrative access to U.S. persons only There will be a separate AWS GovCloud (US) credentials, such as access key and secret access key than the standard AWS account
The user signs in with the IAM user name and password
The AWS GovCloud (US) Region authentication is completely isolated from Amazon.com If the organization is planning to host on EC2 in AWS GovCloud then it will be billed to standard AWS account of organization since AWS GovCloud billing is linked with the standard AWS account and is not be billed separately
http://docs.aws.amazon.com/govcloud-us/latest/UserGuide/whatis.html

QUESTION 180
How does in-memory caching improve the performance of applications in ElastiCache?

A.	It improves application performance by deleting the requests that do not contain frequently accessed data.
B.	It improves application performance by implementing good database indexing strategies.
C.	It improves application performance by using a part of instance RAM for caching important data.
D.	It improves application performance by storing critical pieces of data in memory for low-latency access.

Answer: D
Explanation:
In Amazon ElastiCache, in-memory caching improves application performance by storing critical pieces of data in memory for low-latency access. Cached information may include the results of I/O-intensive database queries or the results of computationally intensive calculations.
http://aws.amazon.com/elasticache/faqs/#g4

QUESTION 181
A user is thinking to use EBS PIOPS volume. Which of the below mentioned options is a right use case for the PIOPS EBS volume?

A.	Analytics
B.	System boot volume
C.	Mongo DB
D.	Log processing

Answer: C
Explanation:
Provisioned IOPS volumes are designed to meet the needs of I/O-intensive workloads, particularly database workloads that are sensitive to storage performance and consistency in random access I/O throughput. Provisioned IOPS volumes are designed to meet the needs of I/O-intensive workloads, particularly database workloads, that are sensitive to storage performance and consistency in random access I/O throughput business applications, database workloads, such as NoSQL DB, RDBMS, etc.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html

QUESTION 182
How can a user list the IAM Role configured as a part of the launch config?

A.	as-describe-launch-configs -iam-profile
B.	as-describe-launch-configs -show-long
C.	as-describe-launch-configs -iam-role
D.	as-describe-launch-configs -role

Answer: B
Explanation:
As-describe-launch-configs describes all the launch config parameters created by the AWS account in the specified region. Generally it returns values, such as Launch Config name, Instance Type and AMI ID. If the user wants additional parameters, such as the IAM Profile used in the config , he has to run command:
as-describe-launch-configs --show-long

QUESTION 183
An organization is setting up a multi-site solution where the application runs on premise as well as on AWS to achieve the minimum recovery time objective(RTO). Which of the below mentioned configurations will not meet the requirements of the multi-site solution scenario?

A.	Configure data replication based on RTO.
B.	Keep an application running on premise as well as in AWS with full capacity.
C.	Setup a single DB instance which will be accessed by both sites.
D.	Setup a weighted DNS service like Route 53 to route traffic across sites.

Answer: C
Explanation:
AWS has many solutions for DR(Disaster recovery) and HA(High Availability). When the organization wants to have HA and DR with multi-site solution, it should setup two sites: one on premise and the other on AWS with full capacity. The organization should setup a weighted DNS service which can route traffic to both sites based on the weightage. When one of the sites fails it can route the entire load to another site. The organization would have minimal RTO in this scenario. If the organization setups a single DB instance, it will not work well in failover. Instead they should have two separate DBs in each site and setup data replication based on RTO(recovery time objective)of the organization.
http://d36cz9buwru1tt.cloudfront.net/AWS_Disaster_Recovery.pdf

QUESTION 184
Which of the following is true of an instance profile when an IAM role is created using the console?

A.	The instance profile uses a different name.
B.	The console gives the instance profile the same name as the role it corresponds to.
C.	The instance profile should be created manually by a user.
D.	The console creates the role and instance profile as separate actions.

Answer: B
Explanation:
Amazon EC2 uses an instance profile as a container for an IAM role. When you create an IAM role using the console, the console creates an instance profile automatically and gives it the same name as the role it corresponds to. If you use the AWS CLI, API, or an AWS SDK to create a role, you create the role and instance profile as separate actions, and you might give them different names.
http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2_instance-profiles.html

QUESTION 185
In the context of policies and permissions in AWS IAM, the Condition element is _____ .

A.	crucial while writing the IAM policies
B.	an optional element
C.	always set to null
D.	a mandatory element

Answer: B
Explanation:
The Condition element (or Condition block) lets you specify conditions for when a policy is in effect. The Condition element is optional.
http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html

QUESTION 186
Which of the following is true while using an IAM role to grant permissions to applications running on Amazon EC2 instances?

A.	All applications on the instance share the same role, but different permissions.
B.	All applications on the instance share multiple roles and permissions.
C.	Multiple roles are assigned to an EC2 instance at a time.
D.	Only one role can be assigned to an EC2 instance at a time.

Answer: D
Explanation:
Only one role can be assigned to an EC2 instance at a time, and all applications on the instance share the same role and permissions.
http://docs.aws.amazon.com/IAM/latest/UserGuide/role-usecase-ec2app.html

QUESTION 187
When using string conditions within IAM, short versions of the available comparators can be used instead of the more verbose ones.
streqi is the short version of the _____ string condition.

A.	StringEqualsIgnoreCase
B.	StringNotEqualsIgnoreCase
C.	StringLikeStringEquals
D.	StringNotEquals

Answer: A
Explanation:
When using string conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions. For instance, streqi is the short version of StringEqualsIgnoreCase that checks for the exact match between two strings ignoring their case.
http://awsdocs.s3.amazonaws.com/SNS/20100331/sns-gsg-2010-03-31.pdf

QUESTION 188
Attempts, one of the three types of items associated with the schedule pipeline in the AWS Data Pipeline, provides robust data management.
Which of the following statements is NOT true about Attempts?

A.	Attempts provide robust data management.
B.	AWS Data Pipeline retries a failed operation until the count of retries reaches the maximum number of allowed retry attempts.
C.	An AWS Data Pipeline Attempt object compiles the pipeline components to create a set of actionable instances.
D.	AWS Data Pipeline Attempt objects track the various attempts, results, and failure reasons if applicable.

Answer: C
Explanation:
Attempts, one of the three types of items associated with a schedule pipeline in AWS Data Pipeline, provides robust data management. AWS Data Pipeline retries a failed operation. It continues to do so until the task reaches the maximum number of allowed retry attempts. Attempt objects track the various attempts, results, and failure reasons if applicable. Essentially, it is the instance with a counter. AWS Data Pipeline performs retries using the same resources from the previous attempts, such as Amazon EMR clusters and EC2 instances.
http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/dp-how-tasks-scheduled.html

QUESTION 189
Select the correct statement about Amazon ElastiCache.

A.	It makes it easy to set up, manage, and scale a distributed in-memory cache environment in the cloud.
B.	It allows you to quickly deploy your cache environment only if you install software.
C.	It does not integrate with other Amazon Web Services.
D.	It cannot run in the Amazon Virtual Private Cloud (Amazon VPC) environment.

Answer: A
Explanation:
ElastiCache is a web service that makes it easy to set up, manage, and scale a distributed in-

memory cache environment in the cloud. It provides a high-performance, scalable, and cost-effective caching solution, while removing the complexity associated with deploying and managing a distributed cache environment. With ElastiCache, you can quickly deploy your cache environment, without having to provision hardware or install software.
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/WhatIs.html

QUESTION 190
In Amazon RDS for PostgreSQL, you can provision up to 3TB storage and 30,000 IOPS per database instance. For a workload with 50% writes and 50% reads running on a cr1.8xlarge instance, you can realize over 25,000 IOPS for PostgreSQL. However, by provisioning more than this limit, you may be able to achieve:

A.	higher latency and lower throughput.
B.	lower latency and higher throughput.
C.	higher throughput only.
D.	higher latency only.

Answer: B
Explanation:
You can provision up to 3TB storage and 30,000 IOPS per database instance. For a workload with 50% writes and 50% reads running on a cr1.8xlarge instance, you can realize over 25,000 IOPS for PostgreSQL. However, by provisioning more than this limit, you may be able to achieve lower latency and higher throughput. Your actual realized IOPS may vary from the amount you provisioned based on your database workload, instance type, and database engine choice.
https://aws.amazon.com/rds/postgresql/

QUESTION 191
Which of the following cannot be done using AWS Data Pipeline?

A.	Create complex data processing workloads that are fault tolerant, repeatable, and highly available.
B.	Regularly access your data where it's stored, transform and process it at scale, and efficiently transfer the results to another AWS service.
C.	Generate reports over data that has been stored.
D.	Move data between different AWS compute and storage services as well as on-premise data sources at specified intervals.

Answer: C
Explanation:
AWS Data Pipeline is a web service that helps you reliably process and move data between different AWS compute and storage services as well as on-premise data sources at specified intervals. With AWS Data Pipeline, you can regularly access your data where it's stored, transform and process it at scale, and efficiently transfer the results to another AWS.
AWS Data Pipeline helps you easily create complex data processing workloads that are fault tolerant, repeatable, and highly available. AWS Data Pipeline also allows you to move and process data that was previously locked up in on-premise data silos.
http://aws.amazon.com/datapipeline/

QUESTION 192
AWS Direct Connect itself has NO specific resources for you to control access to. Therefore, there are no AWS Direct Connect Amazon Resource Names (ARNs) for you to use in an Identity and Access Management (IAM) policy. With that in mind, how is it possible to write a policy to

control access to AWS Direct Connect actions?

A.	You can leave the resource name field blank.
B.	You can choose the name of the AWS Direct Connection as the resource.
C.	You can use an asterisk (*) as the resource.
D.	You can create a name for the resource.

Answer: C
Explanation:
AWS Direct Connect itself has no specific resources for you to control access to. Therefore, there are no AWS Direct Connect ARNs for you to use in an IAM policy. You use an asterisk (*) as the resource when writing a policy to control access to AWS Direct Connect actions.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/using_iam.html

QUESTION 193
Identify an application that polls AWS Data Pipeline for tasks and then performs those tasks.

A.	A task executor
B.	A task deployer
C.	A task runner
D.	A task optimizer

Answer: C
Explanation:
A task runner is an application that polls AWS Data Pipeline for tasks and then performs those tasks. You can either use Task Runner as provided by AWS Data Pipeline, or create a custom Task Runner application.
Task Runner is a default implementation of a task runner that is provided by AWS Data Pipeline. When Task Runner is installed and configured, it polls AWS Data Pipeline for tasks associated with pipelines that you have activated. When a task is assigned to Task Runner, it performs that task and reports its status back to AWS Data Pipeline. If your workflow requires non-default behavior, you'll need to implement that functionality in a custom task runner.
http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/dp-how-remote-taskrunner-client.html

QUESTION 194
With respect to AWS Lambda permissions model, at the time you create a Lambda function, you specify an IAM role that AWS Lambda can assume to execute your Lambda function on your behalf. This role is also referred to as the _____ role.

A.	configuration
B.	execution
C.	delegation
D.	dependency

Answer: B
Explanation:
Regardless of how your Lambda function is invoked, AWS Lambda always executes the function. At the time you create a Lambda function, you specify an IAM role that AWS Lambda can assume to execute your Lambda function on your behalf. This role is also referred to as the execution role.
http://docs.aws.amazon.com/lambda/latest/dg/lambda-dg.pdf

QUESTION 195
Within an IAM policy, can you add an IfExists condition at the end of a Null condition?

A.	Yes, you can add an IfExists condition at the end of a Null condition but not in all Regions.
B.	Yes, you can add an IfExists condition at the end of a Null condition depending on the condition.
C.	No, you cannot add an IfExists condition at the end of a Null condition.
D.	Yes, you can add an IfExists condition at the end of a Null condition.

Answer: C
Explanation:
Within an IAM policy, IfExists can be added to the end of any condition operator except the Null condition. It can be used to indicate that conditional comparison needs to happen if the policy key is present in the context of a request; otherwise, it can be ignored.
http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html

QUESTION 196
Regarding Identity and Access Management (IAM), Which type of special account belonging to your application allows your code to access Google services programmatically?

A.	Service account
B.	Simple Key
C.	OAuth
D.	Code account

Answer: A
Explanation:
A service account is a special Google account that can be used by applications to access Google services programmatically. This account belongs to your application or a virtual machine (VM), instead of to an individual end user. Your application uses the service account to call the Google API of a service, so that the users aren't directly involved.
A service account can have zero or more pairs of service account keys, which are used to authenticate to Google. A service account key is a public/private keypair generated by Google. Google retains the public key, while the user is given the private key.
https://cloud.google.com/iam/docs/service-accounts

QUESTION 197
IAM users do not have permission to create Temporary Security Credentials for federated users and roles by default. In contrast, IAM users can call _____ without the need of any special permissions

A.	GetSessionName
B.	GetFederationToken
C.	GetSessionToken
D.	GetFederationName

Answer: C
Explanation:
Currently the STS API command GetSessionToken is available to every IAM user in your account without previous permission. In contrast, the GetFederationToken command is restricted and explicit permissions need to be granted so a user can issue calls to this particular Action

http://docs.aws.amazon.com/STS/latest/UsingSTS/STSPermission.html

QUESTION 198
An organization is planning to use NoSQL DB for its scalable data needs. The organization wants to host an application securely in AWS VPC. What action can be recommended to the organization?

A.	The organization should setup their own NoSQL cluster on the AWS instance and configure route tables and subnets.
B.	The organization should only use a DynamoDB because by default it is always a part of the default subnet provided by AWS.
C.	The organization should use a DynamoDB while creating a table within the public subnet.
D.	The organization should use a DynamoDB while creating a table within a private subnet.

Answer: A
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Currently VPC does not support DynamoDB. Thus, if the user wants to implement VPC, he has to setup his own NoSQL DB within the VPC. http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Introduction.html

QUESTION 199
What happens when Dedicated instances are launched into a VPC?

A.	If you launch an instance into a VPC that has an instance tenancy of dedicated, you must manually create a Dedicated instance.
B.	If you launch an instance into a VPC that has an instance tenancy of dedicated, your instance is created as a Dedicated instance, only based on the tenancy of the instance.
C.	If you launch an instance into a VPC that has an instance tenancy of dedicated, your instance is automatically a Dedicated instance, regardless of the tenancy of the instance.
D.	None of these are true.

Answer: C
Explanation:
If you launch an instance into a VPC that has an instance tenancy of dedicated, your instance is automatically a Dedicated instance, regardless of the tenancy of the instance.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/dedicated-instance.html

QUESTION 200
An organization is setting up RDS for their applications. The organization wants to secure RDS access with VPC. Which of the following options is not required while designing the RDS with VPC?

A.	The organization must create a subnet group with public and private subnets. Both the subnets can be in the same or separate AZ.
B.	The organization should keep minimum of one IP address in each subnet reserved for RDS failover.
C.	If the organization is connecting RDS from the internet it must enable the VPC attributes DNS hostnames and DNS resolution.
D.	The organization must create a subnet group with VPC using more than one subnet which are a

	part of separate AZs.

Answer: A
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources, such as RDS into a virtual network that the user has defined. Subnets are segments of a VPC's IP address range that the user can designate to a group of VPC resources based on security and operational needs. A DB subnet group is a collection of subnets (generally private) that the user can create in a VPC and assign to the RDS DB instances. A DB subnet group allows the user to specify a particular VPC when creating the DB instances.
Each DB subnet group should have subnets in at least two Availability Zones in a given region. If the RDS instance is required to be accessible from the internet the organization must enable the VPC attributes, DNS hostnames and DNS resolution. For each RDS DB instance that the user runs in a VPC, he should reserve at least one address in each subnet in the DB subnet group for use by Amazon RDS for recovery actions.
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.html

QUESTION 201
You create a VPN connection, and your VPN device supports Border Gateway Protocol (BGP). Which of the following should be specified to configure the VPN connection?

A.	Classless routing
B.	Classfull routing
C.	Dynamic routing
D.	Static routing

Answer: C
Explanation:
If you create a VPN connection, you must specify the type of routing that you plan to use, which will depend upon on the make and model of your VPN devices. If your VPN device supports Border Gateway Protocol (BGP), you need to specify dynamic routing when you configure your VPN connection. If your device does not support BGP, you should specify static routing.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html

QUESTION 202
An organization has developed an application which provides a smarter shopping experience. They need to show a demonstration to various stakeholders who may not be able to access the in premise application so they decide to host a demo version of the application on AWS. Consequently they will need a fixed elastic IP attached automatically to the instance when it is launched. In this scenario which of the below mentioned options will not help assign the elastic IP automatically?

A.	Write a script which will fetch the instance metadata on system boot and assign the public IP using that metadata.
B.	Provide an elastic IP in the user data and setup a bootstrapping script which will fetch that elastic IP and assign it to the instance.
C.	Create a controlling application which launches the instance and assigns the elastic IP based on the parameter provided when that instance is booted.
D.	Launch instance with VPC and assign an elastic IP to the primary network interface.

Answer: A

Explanation:
EC2 allows the user to launch On-Demand instances. If the organization is using an application temporarily only for demo purposes the best way to assign an elastic IP would be:
Launch an instance with a VPC and assign an EIP to the primary network interface. This way on every instance start it will have the same IP
Create a bootstrapping script and provide it some metadata, such as user data which can be used to assign an EIP
Create a controller instance which can schedule the start and stop of the instance and provide an EIP as a parameter so that the controller instance can check the instance boot and assign an EIP
The instance metadata gives the current instance data, such as the public/private IP. It can be of no use for assigning an EIP.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AESDG-chapter-instancedata.html

QUESTION 203
An organization is having a VPC for the HR department, and another VPC for the Admin department. The HR department requires access to all the instances running in the Admin VPC while the Admin department requires access to all the resources in the HR department. How can the organization setup this scenario?

A.	Setup VPC peering between the VPCs of Admin and HR.
B.	Setup ACL with both VPCs which will allow traffic from the CIDR of the other VPC.
C.	Setup the security group with each VPC which allows traffic from the CIDR of another VPC.
D.	It is not possible to connect resources of one VPC from another VPC.

Answer: A
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined.
A VPC peering connection allows the user to route traffic between the peer VPCs using private IP addresses as if they are a part of the same network.
This is helpful when one VPC from the same or different AWS account wants to connect with resources of the other VPC.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-peering.html

QUESTION 204
Can a Direct Connect link be connected directly to the Internet?

A.	Yes, this can be done if you pay for it.
B.	Yes, this can be done only for certain regions.
C.	Yes
D.	No

Answer: D
Explanation:
AWS Direct Connect is a network service that provides an alternative to using the Internet to utilize AWS cloud service. Hence, a Direct Connect link cannot be connected to the Internet directly.
http://aws.amazon.com/directconnect/faqs/

QUESTION 205
ExamKiller has created a multi-tenant Learning Management System (LMS). The application is hosted for five different tenants (clients) in the VPCs of the respective AWS accounts of the

tenant. ExamKiller wants to setup a centralized server which can connect with the LMS of each tenant upgrade if required. ExamKiller also wants to ensure that one tenant VPC should not be able to connect to the other tenant VPC for security reasons. How can ExamKiller setup this scenario?

A.	ExamKiller has to setup one centralized VPC which will peer in to all the other VPCs of the tenants.
B.	ExamKiller should setup VPC peering with all the VPCs peering each other but block the IPs from CIDR of the tenant VPCs to deny them.
C.	ExamKiller should setup all the VPCs with the same CIDR but have a centralized VPC. This way only the centralized VPC can talk to the other VPCs using VPC peering.
D.	ExamKiller should setup all the VPCs meshed together with VPC peering for all VPCs.

Answer: A
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. A VPC peering connection allows the user to route traffic between the peer VPCs using private IP addresses as if they are a part of the same network.
This is helpful when one VPC from the same or different AWS account wants to connect with resources of the other VPC. The organization wants to setup that one VPC can connect with all the other VPCs but all other VPCs cannot connect among each other. This can be achieved by configuring VPC peering where one VPC is peered with all the other VPCs, but the other VPCs are not peered to each other. The VPCs are in the same or a separate AWS account and should not have overlapping CIDR blocks.
http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/peering-configurations-full-access.html# many-vpcs-full-acces

QUESTION 206
A user has created a VPC with CIDR 20.0.0.0/16. The user has created one subnet with CIDR 20.0.0.0/16 in this VPC. The user is trying to create another subnet with the same VPC for CIDR 20.0.0.1/24. What will happen in this scenario?

A.	The VPC will modify the first subnet CIDR automatically to allow the second subnet IP range
B.	The second subnet will be created
C.	It will throw a CIDR overlaps error
D.	It is not possible to create a subnet with the same CIDR as VPC

Answer: C
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. The user can create a subnet with the same size of VPC. However, he cannot create any other subnet since the CIDR of the second subnet will conflict with the first subnet.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

QUESTION 207
True or False: The Amazon ElastiCache clusters are not available for use in VPC at this time.

A.	TRUE
B.	True, but they are available only in the GovCloud.
C.	True, but they are available only on request.

D.	FALSE

Answer: D
Explanation:
Amazon Elasticache clusters can be run in an Amazon VPC. With Amazon VPC, you can define a virtual network topology and customize the network configuration to closely resemble a traditional network that you might operate in your own datacenter. You can now take advantage of the manageability, availability and scalability benefits of Amazon ElastiCache Clusters in your own isolated network. The same functionality of Amazon ElastiCache, including automatic failure detection, recovery, scaling, auto discovery, Amazon CloudWatch metrics, and software patching, are now available in Amazon VPC.
http://aws.amazon.com/about-aws/whats-new/2012/12/20/amazon-elasticache-announces-support-for-amazon-vpc/

QUESTION 208
In Amazon Redshift, how many slices does a dw2.8xlarge node have?

A.	16
B.	8
C.	32
D.	2

Answer: C
Explanation:
The disk storage for a compute node in Amazon Redshift is divided into a number of slices, equal to the number of processor cores on the node. For example, each DW1.XL compute node has two slices, and each DW2.8XL compute node has 32 slices.
http://docs.aws.amazon.com/redshift/latest/dg/t_Distributing_data.html

QUESTION 209
Identify a true statement about using an IAM role to grant permissions to applications running on Amazon EC2 instances.

A.	When AWS credentials are rotated, developers have to update only the root Amazon EC2 instance that uses their credentials.
B.	When AWS credentials are rotated, developers have to update only the Amazon EC2 instance on which the password policy was applied and which uses their credentials.
C.	When AWS credentials are rotated, you don't have to manage credentials and you don't have to worry about long-term security risks.
D.	When AWS credentials are rotated, you must manage credentials and you should consider precautions for long-term security risks.

Answer: C
Explanation:
Using IAM roles to grant permissions to applications that run on EC2 instances requires a bit of extra configuration. Because role credentials are temporary and rotated automatically, you don't have to manage credentials, and you don't have to worry about long-term security risks.
http://docs.aws.amazon.com/IAM/latest/UserGuide/role-usecase-ec2app.html

QUESTION 210
Out of the striping options available for the EBS volumes, which one has the following

disadvantage:
'Doubles the amount of I/O required from the instance to EBS compared to RAID 0, because you're mirroring all writes to a pair of volumes, limiting how much you can stripe.' ?

A.	Raid 1
B.	Raid 0
C.	RAID 1+0 (RAID 10)
D.	Raid 2

Answer: C
Explanation:
RAID 1+0 (RAID 10) doubles the amount of I/O required from the instance to EBS compared to RAID 0, because you're mirroring all writes to a pair of volumes, limiting how much you can stripe. http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html

QUESTION 211
In the context of IAM roles for Amazon EC2, which of the following NOT true about delegating permission to make API requests?

A.	You cannot create an IAM role.
B.	You can have the application retrieve a set of temporary credentials and use them.
C.	You can specify the role when you launch your instances.
D.	You can define which accounts or AWS services can assume the role.

Answer: A
Explanation:
Amazon designed IAM roles so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that the applications use. Instead of creating and distributing your AWS credentials, you can delegate permission to make API requests using IAM roles as follows: Create an IAM role. Define which accounts or AWS services can assume the role. Define which API actions and resources the application can use after assuming the role. Specify the role when you launch your instances. Have the application retrieve a set of temporary credentials and use them.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html

QUESTION 212
In the context of Amazon ElastiCache CLI, which of the following commands can you use to view all ElastiCache instance events for the past 24 hours?

A.	elasticache-events --duration 24
B.	elasticache-events --duration 1440
C.	elasticache-describe-events --duration 24
D.	elasticache describe-events --source-type cache-cluster --duration 1440

Answer: D
Explanation:
In Amazon ElastiCache, the code "aws elasticache describe-events --source-type cache-cluster --duration 1440" is used to list the cache-cluster events for the past 24 hours (1440 minutes).
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/ECEvents.Viewing.html

QUESTION 213

In Amazon Cognito what is a silent push notification?

A.	It is a push message that is received by your application on a user's device that will not be seen by the user.
B.	It is a push message that is received by your application on a user's device that will return the user's geolocation.
C.	It is a push message that is received by your application on a user's device that will not be heard by the user.
D.	It is a push message that is received by your application on a user's device that will return the user's authentication credentials.

Answer: A
Explanation:
Amazon Cognito uses the Amazon Simple Notification Service (SNS) to send silent push notifications to devices. A silent push notification is a push message that is received by your application on a user's device that will not be seen by the user.
http://aws.amazon.com/cognito/faqs/

QUESTION 214
When using Numeric Conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions. Which of the following is the short version of the Numeric Condition "NumericLessThanEquals"?

A.	numlteq
B.	numlteql
C.	numltequals
D.	numeql

Answer: A
Explanation:
When using Numeric Conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions. For instance, numlteq is the short version of NumericLessThanEquals.
http://awsdocs.s3.amazonaws.com/SQS/2011-10-01/sqs-dg-2011-10-01.pdf

QUESTION 215
AWS has launched T2 instances which come with CPU usage credit. An organization has a requirement which keeps an instance running for 24 hours. However, the organization has high usage only during 11 AM to 12 PM. The organization is planning to use a T2 small instance for this purpose. If the organization already has multiple instances running since Jan 2012, which of the below mentioned options should the organization implement while launching a T2 instance?

A.	The organization must migrate to the EC2-VPC platform first before launching a T2 instance.
B.	While launching a T2 instance the organization must create a new AWS account as this account does not have the EC2-VPC platform.
C.	Create a VPC and launch a T2 instance as part of one of the subnets of that VPC.
D.	While launching a T2 instance the organization must select EC2-VPC as the platform.

Answer: C
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. The AWS account provides two

platforms:
EC2-CLASSIC and EC2-VPC, depending on when the user has created his AWS account and which regions he is using. If the user has created the AWS account after 2013-12-04, it supports only EC2-VPC. In this scenario, since the account is before the required date the supported platform will be EC2-CLASSIC. It is required that the organization creates a VPC as the T2 instances can be launched only as a part of VPC.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/vpc-migrate.html

QUESTION 216
How does AWS Data Pipeline execute activities on on-premise resources or AWS resources that you manage?

A.	By supplying a Task Runner package that can be installed on your on-premise hosts
B.	None of these
C.	By supplying a Task Runner file that the resources can access for execution
D.	By supplying a Task Runner json script that can be installed on your on-premise hosts

Answer: A
Explanation:
To enable running activities using on-premise resources, AWS Data Pipeline does the following:
It supply a Task Runner package that can be installed on your on-premise hosts. This package continuously polls the AWS Data Pipeline service for work to perform. When it's time to run a particular activity on your on-premise resources, it will issue the appropriate command to the Task Runner.
https://aws.amazon.com/datapipeline/faqs/

QUESTION 217
Which of following IAM policy elements lets you specify an exception to a list of actions?

A.	NotException
B.	ExceptionAction
C.	Exception
D.	NotAction

Answer: D
Explanation:
The NotAction element lets you specify an exception to a list of actions.
http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html

QUESTION 218
In AWS IAM, which of the following predefined policy condition keys checks how long ago (in seconds) the MFA-validated security credentials making the request were issued using multi-factor authentication (MFA)?

A.	aws:MultiFactorAuthAge
B.	aws:MultiFactorAuthLast
C.	aws:MFAAge
D.	aws:MultiFactorAuthPrevious

Answer: A

Explanation:
aws:MultiFactorAuthAge is one of the predefined keys provided by AWS that can be included within a Condition element of an IAM policy. The key allows to check how long ago (in seconds) the MFA-validated security credentials making the request were issued using Multi-Factor Authentication (MFA).
http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html

QUESTION 219
A user is configuring MySQL RDS with PIOPS. What should be the minimum PIOPS that the user should provision?

A.	1000
B.	200
C.	2000
D.	500

Answer: A
Explanation:
If a user is trying to enable PIOPS with MySQL RDS, the minimum size of storage should be 100 GB and the minimum PIOPS should be 1000.
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.html

QUESTION 220
You are setting up some EBS volumes for a customer who has requested a setup which includes a RAID (redundant array of inexpensive disks). AWS has some recommendations for RAID setups. Which RAID setup is not recommended for Amazon EBS?

A.	RAID 1 only
B.	RAID 5 only
C.	RAID 5 and RAID 6
D.	RAID 0 only

Answer: C
Explanation:
With Amazon EBS, you can use any of the standard RAID configurations that you can use with a traditional bare metal server, as long as that particular RAID configuration is supported by the operating system for your instance. This is because all RAID is accomplished at the software level. For greater I/O performance than you can achieve with a single volume, RAID 0 can stripe multiple volumes together; for on-instance redundancy, RAID 1 can mirror two volumes together. RAID 5 and RAID 6 are not recommended for Amazon EBS because the parity write operations of these RAID modes consume some of the IOPS available to your volumes.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html

QUESTION 221
Once the user has set ElastiCache for an application and it is up and running, which services, does Amazon not provide for the user:

A.	The ability for client programs to automatically identify all of the nodes in a cache cluster, and to initiate and maintain connections to all of these nodes
B.	Automating common administrative tasks such as failure detection and recovery, and software

	patching
C.	Providing default Time To Live (TTL) in the AWS Elasticache Redis Implementation for different type of data.
D.	Providing detailed monitoring metrics associated with your Cache Nodes, enabling you to diagnose and react to issues very quickly

Answer: C
Explanation:
Amazon provides failure detection and recovery, and software patching and monitoring tools which is called CloudWatch. In addition it provides also Auto Discovery to automatically identify and initialize all nodes of cache cluster for Amazon ElastiCache.
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/WhatIs.html

QUESTION 222
In the context of AWS Cloud Hardware Security Module(HSM), does your application need to reside in the same VPC as the CloudHSM instance?

A.	No, but the server or instance on which your application and the HSM client is running must have network (IP) reachability to the HSM.
B.	Yes, always
C.	No, but they must reside in the same Availability Zone.
D.	No, but it should reside in same Availability Zone as the DB instance.

Answer: A
Explanation:
Your application does not need to reside in the same VPC as the CloudHSM instance. However, the server or instance on which your application and the HSM client is running must have network (IP) reachability to the HSM. You can establish network connectivity in a variety of ways, including operating your application in the same VPC, with VPC peering, with a VPN connection, or with Direct Connect.
https://aws.amazon.com/cloudhsm/faqs/

QUESTION 223
True or False: In Amazon ElastiCache, you can use Cache Security Groups to configure the cache clusters that are part of a VPC.

A.	FALSE
B.	TRUE
C.	True, this is applicable only to cache clusters that are running in an Amazon VPC environment.
D.	True, but only when you configure the cache clusters using the Cache Security Groups from the console navigation pane.

Answer: A
Explanation:
Amazon ElastiCache cache security groups are only applicable to cache clusters that are not running in an Amazon Virtual Private Cloud environment (VPC). If you are running in an Amazon Virtual Private Cloud, Cache Security Groups is not available in the console navigation pane.
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/CacheSecurityGroup.html

QUESTION 224
What is the role of the PollForTask action when it is called by a task runner in AWS Data

Pipeline?

A.	It is used to retrieve the pipeline definition.
B.	It is used to report the progress of the task runner to AWS Data Pipeline.
C.	It is used to receive a task to perform from AWS Data Pipeline.
D.	It is used to inform AWS Data Pipeline of the outcome when the task runner completes a task.

Answer: C
Explanation:
Task runners call PollForTask to receive a task to perform from AWS Data Pipeline. If tasks are ready in the work queue, PollForTask returns a response immediately. If no tasks are available in the queue, PollForTask uses long-polling and holds on to a poll connection for up to 90 seconds, during which time any newly scheduled tasks are handed to the task agent. Your remote worker should not call PollForTask again on the same worker group until it receives a response, and this may take up to 90 seconds.
http://docs.aws.amazon.com/datapipeline/latest/APIReference/API_PollForTask.html

QUESTION 225
What is the average queue length recommended by AWS to achieve a lower latency for the 200 PIOPS EBS volume?

A.	5
B.	1
C.	2
D.	4

Answer: B
Explanation:
The queue length is the number of pending I/O requests for a device. The optimal average queue length will vary for every customer workload, and this value depends on a particular application's sensitivity to IOPS and latency. If the workload is not delivering enough I/O requests to maintain the optimal average queue length, then the EBS volume might not consistently deliver the IOPS that have been provisioned. However, if the workload maintains an average queue length that is higher than the optimal value, then the per-request I/O latency will increase; in this case, the user should provision more IOPS for his volume. AWS recommends that the user should target an optimal average queue length of 1 for every 200 provisioned IOPS and tune that value based on his application requirements.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-workload-demand.html

QUESTION 226
Who is responsible for modifying the routing tables and networking ACLs in a VPC to ensure that a DB instance is reachable from other instances in the VPC?

A.	AWS administrators
B.	The owner of the AWS account
C.	Amazon
D.	The DB engine vendor

Answer: B
Explanation:
You are in charge of configuring the routing tables of your VPC as well as the network ACLs rules needed to make your DB instances accessible from all the instances of your VPC that need to

communicate with it.
http://aws.amazon.com/rds/faqs/

QUESTION 227
An organization is planning to host a web application in the AWS VPC. The organization does not want to host a database in the public cloud due to statutory requirements. How can the organization setup in this scenario?

A.	The organization should plan the app server on the public subnet and database in the organization's data center and connect them with the VPN gateway.
B.	The organization should plan the app server on the public subnet and use RDS with the private subnet for a secure data operation.
C.	The organization should use the public subnet for the app server and use RDS with a storage gateway to access as well as sync the data securely from the local data center.
D.	The organization should plan the app server on the public subnet and database in a private subnet so it will not be in the public cloud.

Answer: A
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. When the user has configured this setup with Wizard, it will create a virtual private gateway to route all the traffic of the VPN subnet. If the virtual private gateway is attached with VPC and the user deletes the VPC from the console it will first automatically detach the gateway and only then delete the VPC.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

QUESTION 228
A user is trying to create a PIOPS EBS volume with 4000 IOPS and 100 GB size. AWS does not allow the user to create this volume. What is the possible root cause for this?

A.	PIOPS is supported for EBS higher than 500 GB size
B.	The maximum IOPS supported by EBS is 3000
C.	The ratio between IOPS and the EBS volume is higher than 30
D.	The ratio between IOPS and the EBS volume is lower than 50

Answer: C
Explanation:
A Provisioned IOPS (SSD) volume can range in size from 4 GiB to 16 TiB and you can provision up to 20,000 IOPS per volume. The ratio of IOPS provisioned to the volume size requested should be a maximum of 30; for example, a volume with 3000 IOPS must be at least 100 GB.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html#EBSVolumeTypes_piops

QUESTION 229
A user is planning to host a Highly Available system on the AWS VPC. Which of the below mentioned statements is helpful in this scenario?

A.	Create VPC subnets in two separate availability zones and launch instances in different subnets.
B.	Create VPC with only one public subnet and launch instances in different AZs using that subnet.

C.	Create two VPCs in two separate zones and setup failover with ELB such that if one VPC fails it will divert traffic to another VPC.
D.	Create VPC with only one private subnet and launch instances in different AZs using that subnet.

Answer: A
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. The VPC is always specific to a region. The user can create a VPC which can span multiple Availability Zones by adding one or more subnets in each Availability Zone. Each subnet must reside entirely within one Availability Zone and cannot span across zones.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html#VPCSubnet

QUESTION 230
A user is creating a PIOPS volume. What is the maximum ratio the user should configure between PIOPS and the volume size?

A.	5
B.	10
C.	20
D.	30

Answer: D
Explanation:
Provisioned IOPS volumes are designed to meet the needs of I/O-intensive workloads, particularly database workloads that are sensitive to storage performance and consistency in random access I/O throughput. A provisioned IOPS volume can range in size from 10 GB to 1 TB and the user can provision up to 4000 IOPS per volume.
The ratio of IOPS provisioned to the volume size requested can be a maximum of 30; for example, a volume with 3000 IOPS must be at least 100 GB.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html

QUESTION 231
What is a possible reason you would need to edit claims issued in a SAML token?

A.	The NameIdentifier claim cannot be the same as the username stored in AD.
B.	Authentication fails consistently.
C.	The NameIdentifier claim cannot be the same as the claim URI.
D.	The NameIdentifier claim must be the same as the username stored in AD.

Answer: A
Explanation:
The two reasons you would need to edit claims issued in a SAML token are:
The NameIdentifier claim cannot be the same as the username stored in AD, and The app requires a different set of claim URIs.
https://azure.microsoft.com/en-us/documentation/articles/active-directory-saml-claims-customization/

QUESTION 232
A government client needs you to set up secure cryptographic key storage for some of their extremely confidential data. You decide that the AWS CloudHSM is the best service for this.

However, there seem to be a few pre-requisites before this can happen, one of those being a security group that has certain ports open. Which of the following is correct in regards to those security groups?

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

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

QUESTION 233
What is the network performance offered by the c4.8xlarge instance in Amazon EC2?

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

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

QUESTION 234
An organization is setting up a web application with the JEE stack. The application uses the JBoss app server and MySQL DB. The application has a logging module which logs all the activities whenever a business function of the JEE application is called. The logging activity takes some time due to the large size of the log file. If the application wants to setup a scalable infrastructure which of the below mentioned options will help achieve this setup?

A.	Host the log files on EBS with PIOPS which will have higher I/O.
B.	Host logging and the app server on separate servers such that they are both in the same zone.
C.	Host logging and the app server on the same instance so that the network latency will be shorter.
D.	Create a separate module for logging and using SQS compartmentalize the module such that all calls to logging are asynchronous.

Answer: D
Explanation:
The organization can always launch multiple EC2 instances in the same region across multiple AZs for HA and DR. The AWS architecture practice recommends compartmentalizing the functionality such that they can both run in parallel without affecting the performance of the main application. In this scenario logging takes a longer time due to the large size of the log file. Thus, it is recommended that the organization should separate them out and make separate modules and make asynchronous calls among them. This way the application can scale as per the requirement and the performance will not bear the impact of logging.
http://www.awsarchitectureblog.com/2014/03/aws-and-compartmentalization.html

QUESTION 235
You're trying to delete an SSL certificate from the IAM certificate store, and you're getting the message "Certificate: <certificate-id> is being used by CloudFront." Which of the following statements is probably the reason why you are getting this error?

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

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

QUESTION 236
A user has set the IAM policy where it denies all requests if a request is not from IP 10.10.10.1/32. The other policy says allow all requests between 5 PM to 7 PM. What will happen when a user is requesting access from IP 55.109.10.12/32 at 6 PM?

A.	It will deny access
B.	It is not possible to set a policy based on the time or IP
C.	IAM will throw an error for policy conflict
D.	It will allow access

Answer: A
Explanation:
When a request is made, the AWS IAM policy decides whether a given request should be allowed or denied. The evaluation logic follows these rules:
By default, all requests are denied. (In general, requests made using the account credentials for resources in the account are always allowed.)
An explicit allow policy overrides this default.
An explicit deny policy overrides any allows.
In this case since there are explicit deny and explicit allow statements. Thus, the request will be denied since deny overrides allow.

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

QUESTION 237
Do you need to use Amazon Cognito to use the Amazon Mobile Analytics service?

A.	No. However, it is recommend by AWS to use Amazon Cognito for security best practices.
B.	Yes. You need to use it only if you have IAM root access.
C.	No. You cannot use it at all, and you need to use AWS IAM accounts.
D.	Yes. It is recommended by AWS to use Amazon Cognito to use Amazon Mobile Analytics service.

Answer: A
Explanation:
You can initialize Amazon Mobile Analytics using AWS IAM accounts. AWS recommend using Amazon Cognito for security best practices.
http://aws.amazon.com/mobileanalytics/faqs/

QUESTION 238
Which of the following AWS services can be used to define alarms to trigger on a certain activity, such as activity success, failure, or delay in AWS Data Pipeline?

A.	Amazon SES
B.	Amazon CodeDeploy
C.	Amazon SNS
D.	Amazon SQS

Answer: C
Explanation:
In AWS Data Pipeline, you can define Amazon SNS alarms to trigger on activities such as success, failure, or delay by creating an alarm object and referencing it in the onFail, onSuccess, or onLate slots of the activity object.
https://aws.amazon.com/datapipeline/faqs/

QUESTION 239
You want to use Amazon Redshift and you are planning to deploy dw1.8xlarge nodes. What is the minimum amount of nodes that you need to deploy with this kind of configuration?

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

Answer: D
Explanation:
For a single-node configuration in Amazon Redshift, the only option available is the smallest of the two options. The 8XL extra-large nodes are only available in a multi-node configuration
http://docs.aws.amazon.com/redshift/latest/mgmt/working-with-clusters.html

QUESTION 240
Mike is appointed as Cloud Consultant in ExamKiller.com. ExamKiller has the following VPCs set-up in the US East Region:

A VPC with CIDR block 10.10.0.0/16, a subnet in that VPC with CIDR block 10.10.1.0/24 A VPC with CIDR block 10.40.0.0/16, a subnet in that VPC with CIDR block 10.40.1.0/24 ExamKiller.com is trying to establish network connection between two subnets, a subnet with CIDR block 10.10.1.0/24 and another subnet with CIDR block 10.40.1.0/24. Which one of the following solutions should Mike recommend to ExamKiller.com?

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

Answer: C
Explanation:
A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IP addresses. EC2 instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, or with a VPC in another AWS account within a single region. AWS uses the existing infrastructure of a VPC to create a VPC peering connection; it is neither a gateway nor a VPN connection, and does not rely on a separate piece of physical hardware.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-peering.html

QUESTION 241
Can Provisioned IOPS be used on RDS instances launched in a VPC?

A.	Yes, they can be used only with Oracle based instances.
B.	Yes, they can be used for all RDS instances.
C.	No
D.	Yes, they can be used only with MySQL based instances.

Answer: B
Explanation:
The basic building block of Amazon RDS is the DB instance. DB instance storage comes in three types:
Magnetic, General Purpose (SSD), and Provisioned IOPS (SSD). When you buy a server, you get CPU, memory, storage, and IOPS, all bundled together. With Amazon RDS, these are split apart so that you can scale them independently. So, for example, if you need more CPU, less IOPS, or more storage, you can easily allocate them.
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/RDSFAQ.PIOPS.html

QUESTION 242
To get started using AWS Direct Connect, in which of the following steps do you configure Border Gateway Protocol (BGP)?

A.	Complete the Cross Connect
B.	Configure Redundant Connections with AWS Direct Connect
C.	Create a Virtual Interface
D.	Download Router Configuration

Answer: C
Explanation:

In AWS Direct Connect, your network must support Border Gateway Protocol (BGP) and BGP MD5 authentication, and you need to provide a private Autonomous System Number (ASN) for that to connect to Amazon Virtual Private Cloud (VPC). To connect to public AWS products such as Amazon EC2 and Amazon S3, you will also need to provide a public ASN that you own (preferred) or a private ASN. You have to configure BGP in the Create a Virtual Interface step.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/getstarted.html#createvirtualinterface

QUESTION 243
Which of the following components of AWS Data Pipeline polls for tasks and then performs those tasks?

A.	Pipeline Definition
B.	Task Runner
C.	Amazon Elastic MapReduce (EMR)
D.	AWS Direct Connect

Answer: B
Explanation:
Task Runner polls for tasks and then performs those tasks.
http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/what-is-datapipeline.html

QUESTION 244
A user is hosting a public website on AWS. The user wants to have the database and the app server on the AWS VPC. The user wants to setup a database that can connect to the Internet for any patch upgrade but cannot receive any request from the internet. How can the user set this up?

A.	Setup DB in a private subnet with the security group allowing only outbound traffic.
B.	Setup DB in a public subnet with the security group allowing only inbound data.
C.	Setup DB in a local data center and use a private gateway to connect the application with DB.
D.	Setup DB in a private subnet which is connected to the internet via NAT for outbound.

Answer: D
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. AWS provides two features that the user can use to increase security in VPC: security groups and network ACLs. When the user wants to setup both the DB and App on VPC, the user should make one public and one private subnet. The DB should be hosted in a private subnet and instances in that subnet cannot reach the internet. The user can allow an instance in his VPC to initiate outbound connections to the internet but prevent unsolicited inbound connections from the internet by using a Network Address Translation (NAT) instance.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

QUESTION 245
An organization is setting up their website on AWS. The organization is working on various security measures to be performed on the AWS EC2 instances. Which of the below mentioned security mechanisms will not help the organization to avoid future data leaks and identify security weaknesses?

A.	Run penetration testing on AWS with prior approval from Amazon.
B.	Perform SQL injection for application testing.
C.	Perform a Code Check for any memory leaks.
D.	Perform a hardening test on the AWS instance.

Answer: C
Explanation:
AWS security follows the shared security model where the user is as much responsible as Amazon. Since Amazon is a public cloud it is bound to be targeted by hackers. If an organization is planning to host their application on AWS EC2, they should perform the below mentioned security checks as a measure to find any security weakness/data leaks:
Perform penetration testing as performed by attackers to find any vulnerability. The organization must take an approval from AWS before performing penetration testing Perform hardening testing to find if there are any unnecessary ports open Perform SQL injection to find any DB security issues
The code memory checks are generally useful when the organization wants to improve the application performance.
http://aws.amazon.com/security/penetration-testing/

QUESTION 246
In Amazon ElastiCache, the default cache port is:

A.	for Memcached 11210 and for Redis 6380.
B.	for Memcached 11211 and for Redis 6380.
C.	for Memcached 11210 and for Redis 6379.
D.	for Memcached 11211 and for Redis 6379.

Answer: D
Explanation:
In Amazon ElastiCache, you can specify a new port number for your cache cluster, which by default is 11211 for Memcached and 6379 for Redis.
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/GettingStarted.AuthorizeAccess.html

QUESTION 247
A user has created a VPC with public and private subnets using the VPC wizard. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.0.0/24 . The NAT instance ID is i-a12345. Which of the below mentioned entries are required in the main route table attached with the private subnet to allow instances to connect with the internet?

A.	Destination: 20.0.0.0/0 and Target: 80
B.	Destination: 20.0.0.0/0 and Target: i-a12345
C.	Destination: 20.0.0.0/24 and Target: i-a12345
D.	Destination: 0.0.0.0/0 and Target: i-a12345

Answer: D
Explanation:
A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the Internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create two route tables and attach to the subnets. The main route table will have the entry "Destination: 0.0.0.0/0 and Target: i-a12345", which allows all the

instances in the private subnet to connect to the internet using NAT.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

QUESTION 248
Which of the following cannot be used to manage Amazon ElastiCache and perform administrative tasks?

A.	AWS software development kits (SDKs)
B.	Amazon S3
C.	ElastiCache command line interface (CLI)
D.	AWS CloudWatch

Answer: D
Explanation:
CloudWatch is a monitoring tool and doesn't give users access to manage Amazon ElastiCache.
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/WhatIs.Managing.html

QUESTION 249
Which of the following statements is correct about AWS Direct Connect?

A.	Connections to AWS Direct Connect require double clad fiber for 1 gigabit Ethernet with Auto Negotiation enabled for the port.
B.	An AWS Direct Connect location provides access to Amazon Web Services in the region it is associated with.
C.	AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 50 gigabit Ethernet cable.
D.	To use AWS Direct Connect, your network must be colocated with a new AWS Direct Connect location.

Answer: B
Explanation:
AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 1 gigabit or 10 gigabit Ethernet fiber-optic cable. An AWS Direct Connect location provides access to Amazon Web Services in the region it is associated with, as well as access to other US regions. To use AWS Direct Connect, your network is colocated with an existing AWS Direct Connect location. Connections to AWS Direct Connect require single mode fiber, 1000BASE-LX (1310nm) for 1 gigabit Ethernet, or 10GBASE-LR (1310nm) for 10 gigabit Ethernet. Auto Negotiation for the port must be disabled.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

QUESTION 250
Which of the following statements is correct about the number of security groups and rules applicable for an EC2-Classic instance and an EC2-VPC network interface?

A.	In EC2-Classic, you can associate an instance with up to 5 security groups and add up to 50 rules to a security group. In EC2-VPC, you can associate a network interface with up to 500 security groups and add up to 100 rules to a security group.
B.	In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 50 rules to a security group. In EC2-VPC, you can associate a network interface with up to 5 security groups and add up to 100 rules to a security group.
C.	In EC2-Classic, you can associate an instance with up to 5 security groups and add up to 100

	rules to a security group. In EC2-VPC, you can associate a network interface with up to 500 security groups and add up to 50 rules to a security group.
D.	In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 100 rules to a security group. In EC2-VPC, you can associate a network interface with up to 5 security groups and add up to 50 rules to a security group.

Answer: D
Explanation:
A security group acts as a virtual firewall that controls the traffic for one or more instances. When you launch an instance, you associate one or more security groups with the instance. You add rules to each security group that allow traffic to or from its associated instances. If you're using EC2-Classic, you must use security groups created specifically for EC2-Classic. In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 100 rules to a security group. If you're using EC2-VPC, you must use security groups created specifically for your VPC. In EC2-VPC, you can associate a network interface with up to 5 security groups and add up to 50 rules to a security group.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html

QUESTION 251
An organization is hosting a scalable web application using AWS. The organization has configured internet facing ELB and Auto Scaling to make the application scalable. Which of the below mentioned statements is required to be followed when the application is planning to host a web application on VPC?

A.	The ELB can be in a public or a private subnet but should have the ENI which is attached to an elastic IP.
B.	The ELB must not be in any subnet; instead it should face the internet directly.
C.	The ELB must be in a public subnet of the VPC to face the internet traffic.
D.	The ELB can be in a public or a private subnet but must have routing tables attached to divert the internet traffic to it.

Answer: C
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances. There are two ELBs available with VPC:
internet facing and internal (private) ELB. For internet facing ELB it is required that ELB should be in a public subnet.
After the user creates the public subnet, he should ensure to associate the route table of the public subnet with the internet gateway to enable the load balancer in the subnet to connect with the internet.
http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/CreateVPCForELB.html

QUESTION 252
Is there any way to own a direct connection to Amazon Web Services?

A.	No, AWS only allows access from the public Internet.
B.	No, you can create an encrypted tunnel to VPC, but you cannot own the connection.
C.	Yes, you can via Amazon Dedicated Connection.

D.	Yes, you can via AWS Direct Connect.

Answer: D
Explanation:
AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 1 gigabit or 10 gigabit Ethernet fiber-optic cable. One end of the cable is connected to your router, the other to an AWS Direct Connect router. With this connection in place, you can create virtual interfaces directly to the AWS cloud (for example, to Amazon Elastic Compute Cloud (Amazon EC2) and Amazon Simple Storage Service (Amazon S3)) and to Amazon Virtual Private Cloud (Amazon VPC), bypassing Internet service providers in your network path.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

QUESTION 253
Identify a true statement about the statement ID (Sid) in IAM.

A.	You cannot expose the Sid in the IAM API.
B.	You cannot use a Sid value as a sub-ID for a policy document's ID for services provided by SQS and SNS.
C.	You can expose the Sid in the IAM API.
D.	You cannot assign a Sid value to each statement in a statement array.

Answer: A
Explanation:
The Sid(statement ID) is an optional identifier that you provide for the policy statement. You can assign a Sid a value to each statement in a statement array. In IAM, the Sid is not exposed in the IAM API. You can't retrieve a particular statement based on this ID.
http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html#Sid

QUESTION 254
In Amazon ElastiCache, which of the following statements is correct?

A.	When you launch an ElastiCache cluster into an Amazon VPC private subnet, every cache node is assigned a public IP address within that subnet.
B.	You cannot use ElastiCache in a VPC that is configured for dedicated instance tenancy.
C.	If your AWS account supports only the EC2-VPC platform, ElastiCache will never launch your cluster in a VPC.
D.	ElastiCache is not fully integrated with Amazon Virtual Private Cloud (VPC).

Answer: B
Explanation:
The VPC must allow non-dedicated EC2 instances. You cannot use ElastiCache in a VPC that is configured for dedicated instance tenancy.
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/AmazonVPC.EC.html

QUESTION 255
An organization has setup RDS with VPC. The organization wants RDS to be accessible from the internet. Which of the below mentioned configurations is not required in this scenario?

A.	The organization must enable the parameter in the console which makes the RDS instance publicly accessible.
B.	The organization must allow access from the internet in the RDS VPC security group.

C.	The organization must setup RDS with the subnet group which has an external IP.
D.	The organization must enable the VPC attributes DNS hostnames and DNS resolution.

Answer: C
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources, such as RDS into a virtual network that the user has defined. Subnets are segments of a VPC's IP address range that the user can designate to a group of VPC resources based on security and operational needs. A DB subnet group is a collection of subnets (generally private) that the user can create in a VPC and which the user assigns to the RDS DB instances. A DB subnet group allows the user to specify a particular VPC when creating DB instances. If the RDS instance is required to be accessible from the internet:
The organization must setup that the RDS instance is enabled with the VPC attributes, DNS hostnames and DNS resolution.
The organization must enable the parameter in the console which makes the RDS instance publicly accessible.
The organization must allow access from the internet in the RDS VPC security group.
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.html

QUESTION 256
An organization, which has the AWS account ID as 999988887777, has created 50 IAM users. All the users are added to the same group examkiller. If the organization has enabled that each IAM user can login with the AWS console, which AWS login URL will the IAM users use??

A.	https://999988887777.aws.amazon.com/examkiller/
B.	https://signin.aws.amazon.com/examkiller/
C.	https://examkiller.signin.aws.amazon.com/999988887777/console/
D.	https://999988887777.signin.aws.amazon.com/console/

Answer: D
Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. Once the organization has created the IAM users, they will have a separate AWS console URL to login to the AWS console. The console login URL for the IAM user will be https:// AWS_Account_ID.signin.aws.amazon.com/console/. It uses only the AWS account ID and does not depend on the group or user ID.
http://docs.aws.amazon.com/IAM/latest/UserGuide/AccountAlias.html

QUESTION 257
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 AWS 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 SAML 2.0-compliam identity provider (IDP) to retrieve temporary security

	credentials to enable NOC members to sign in to the AWS Management Console.

Answer: C

QUESTION 258
You are looking to migrate your Development (Dev) and Test environments to AWS. You have decided to use separate AWS accounts to host each environment. You plan to link each accounts bill to a Master AWS account using Consolidated Billing. To make sure you Keep within budget you would like to implement a way for administrators in the Master account to have access to stop, delete and/or terminate resources in both the Dev and Test accounts. Identify which option will allow you to achieve this goal.

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

Answer: C

QUESTION 259
Your company hosts a social media website for storing and sharing documents. The web application allows user to upload large files while resuming and pausing the upload as needed. Currently, files are uploaded to your PHP front end backed by Elastic load Balancing and an autoscaling fleet of Amazon Elastic Compute Cloud (EC2) instances that scale upon average of bytes received (NetworkIn). After a file has been uploaded, it is copied to Amazon Simple Storage Service (S3). Amazon EC2 instances use an AWS Identity and Access Management (IAM) role that allows Amazon S3 uploads. Over the last six months, your user base and scale have increased significantly, forcing you to increase the Auto Scaling group's Max parameter a few times. Your CFO is concerned about rising costs and has asked you to adjust the architecture where needed to better optimize costs. Which architecture change could you introduce to reduce costs and still keep your web application secure and scalable?

A.	Replace the Auto Scaling launch configuration to include c3.8xlarge instances; those instances can potentially yield a network throuthput of 10gbps.
B.	Re-architect your ingest pattern, have the app authenticate against your identity provider, and use your identity provider as a broker fetching temporary AWS credentials from AWS Secure Token Service (GetFederationToken). Securely pass the credentials and S3 endpoint/prefix to your app. Implement client-side logic to directly upload the file to Amazon S3 using the given credentials and S3 prefix.
C.	Re-architect your ingest pattern, and move your web application instances into a VPC public subnet. Attach a public IP address for each EC2 instance (using the Auto Scaling launch configuration settings). Use Amazon Route 53 Round Robin records set and HTTP health check to DNS load balance the app requests; this approach will significantly reduce the cost by bypassing Elastic Load Balancing.
D.	Re-architect your ingest pattern, have the app authenticate against your identity provider, and use

your identity provider as a broker fetching temporary AWS credentials from AWS Secure Token Service (GetFederationToken). Securely pass the credentials and S3 endpoint/prefix to your app. Implement client-side logic that used the S3 multipart upload API to directly upload the file to Amazon S3 using the given credentials and S3 prefix.

Answer: C

QUESTION 260
You are designing a personal document-archiving solution for your global enterprise with thousands of employee. Each employee has potentially gigabytes of data to be backed up in this archiving solution. The solution will be exposed to the employees as an application, where they can just drag and drop their files to the archiving system. Employees can retrieve their archives through a web interface. The corporate network has high bandwidth AWS Direct Connect connectivity to AWS. You have a regulatory requirement that all data needs to be encrypted before being uploaded to the cloud.
How do you implement this in a highly available and cost-efficient way?

A.	Manage encryption keys on-premises in an encrypted relational database. Set up an on-premises server with sufficient storage to temporarily store files, and then upload them to Amazon S3, providing a client-side master key.
B.	Mange encryption keys in a Hardware Security Module (HSM) appliance on-premises serve r with sufficient storage to temporarily store, encrypt, and upload files directly into Amazon Glacier.
C.	Manage encryption keys in Amazon Key Management Service (KMS), upload to Amazon Simple Storage Service (S3) with client-side encryption using a KMS customer master key ID, and configure Amazon S3 lifecycle policies to store each object using the Amazon Glacier storage tier.
D.	Manage encryption keys in an AWS CloudHSM appliance. Encrypt files prior to uploading on the employee desktop, and then upload directly into Amazon Glacier.

Answer: C

QUESTION 261
You are responsible for a web application that consists of an Elastic Load Balancing (ELB) load balancer in front of an Auto Scaling group of Amazon Elastic Compute Cloud (EC2) instances. For a recent deployment of a new version of the application, a new Amazon Machine Image (AMI) was created, and the Auto Scaling group was updated with a new launch configuration that refers to this new AMI. During the deployment, you received complaints from users that the website was responding with errors. All instances passed the ELB health checks.
What should you do in order to avoid errors for future deployments? (Choose 2 answer)

A.	Add an Elastic Load Balancing health check to the Auto Scaling group. Set a short period for the health checks to operate as soon as possible in order to prevent premature registration of the instance to the load balancer.
B.	Enable EC2 instance CloudWatch alerts to change the launch configuration's AMI to the previous one. Gradually terminate instances that are using the new AMI.
C.	Set the Elastic Load Balancing health check configuration to target a part of the application that fully tests application health and returns an error if the tests fail.
D.	Create a new launch configuration that refers to the new AMI, and associate it with the group. Double the size of the group, wait for the new instances to become healthy, and reduce back to the original size. If new instances do not become healthy, associate the previous launch configuration.

E.	Increase the Elastic Load Balancing Unhealthy Threshold to a higher value to prevent an unhealthy instance from going into service behind the load balancer.

Answer: CD

QUESTION 262
Which is a valid Amazon Resource name (ARN) for IAM?

A.	aws:iam::123456789012:instance-profile/Webserver
B.	arn:aws:iam::123456789012:instance-profile/Webserver
C.	123456789012:aws:iam::instance-profile/Webserver
D.	arn:aws:iam::123456789012::instance-profile/Webserver

Answer: B

QUESTION 263
Dave is the main administrator in Example Corp., and he decides to use paths to help delineate the users in the company and set up a separate administrator group for each path-based division. Following is a subset of the full list of paths he plans to use:

```
. /marketing
. /sales
. /legal
```

Dave creates an administrator group for the marketing part of the company and calls it Marketing_Admin.
He assigns it the /marketing path. The group's ARN is arn:aws:iam::123456789012:group/marketing/Marketing_Admin. Dave assigns the following policy to the Marketing_Admin group that gives the group permission to use all IAM actions with all groups and users in the /marketing path. The policy also gives the Marketing_Admin group permission to perform any AWS S3 actions on the objects in the portion of the corporate bucket.

```
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Deny",
"Action": "iam:*",
"Resource": [
"arn:aws:iam::123456789012:group/marketing/*",
"arn:aws:iam::123456789012:user/marketing/*"
]
},
{
"Effect": "Allow",
"Action": "s3:*",
"Resource": "arn:aws:s3:::example_bucket/marketing/*"
},
{
"Effect": "Allow",
"Action": "s3:ListBucket*",
"Resource": "arn:aws:s3:::example_bucket",
"Condition":{"StringLike":{"s3:prefix": "marketing/*"}}
```

```
}
]
}
```

A.	True
B.	False

Answer: B

QUESTION 264
How can an EBS volume that is currently attached to an EC2 instance be migrated from one Availability Zone to another?

A.	Detach the volume and attach it to another EC2 instance in the other AZ.
B.	Simply create a new volume in the other AZ and specify the original volume as the source.
C.	Create a snapshot of the volume, and create a new volume from the snapshot in the other AZ.
D.	Detach the volume, then use the ec2-migrate-voiume command to move it to another AZ.

Answer: C

QUESTION 265
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

QUESTION 266
Which of the following are characteristics of Amazon VPC subnets? Choose 2 answers

A.	Each subnet spans at least 2 Availability Zones to provide a high-availability environment.
B.	Each subnet maps to a single Availability Zone.
C.	CIDR block mask of /25 is the smallest range supported.
D.	By default, all subnets can route between each other, whether they are private or public.
E.	Instances in a private subnet can communicate with the Internet only if they have an Elastic IP.

Answer: AE

QUESTION 267
In AWS, which security aspects are the customer's responsibility? Choose 4 answers

A.	Security Group and ACL (Access Control List) settings
B.	Decommissioning storage devices
C.	Patch management on the EC2 instance's operating system
D.	Life-cycle management of IAM credentials
E.	Controlling physical access to compute resources
F.	Encryption of EBS (Elastic Block Storage) volumes

Answer: ACDF

QUESTION 268
When you put objects in Amazon S3, what is the indication that an object was successfully stored?

A.	A HTTP 200 result code and MD5 checksum, taken together, indicate that the operation was successful.
B.	Amazon S3 is engineered for 99.999999999% durability. Therefore there is no need to confirm that data was inserted.
C.	A success code is inserted into the S3 object metadata.
D.	Each S3 account has a special bucket named _s3_logs. Success codes are written to this bucket witha timestamp and checksum.

Answer: A

QUESTION 269
A customer is deploying an SSL enabled web application to AWS and would like to implement a separation of roles between the EC2 service administrators that are entitled to login to instances as well as making API calls and the security officers who will maintain and have exclusive access to the application's X.509 certificate that contains the private key.

A.	Upload the certificate on an S3 bucket owned by the security officers and accessible only by EC2 Role of the web servers.
B.	Configure the web servers to retrieve the certificate upon boot from an CloudHSM is managed by the security officers.
C.	Configure system permissions on the web servers to restrict access to the certificate only to the authority security officers
D.	Configure IAM policies authorizing access to the certificate store only to the security officers and terminate SSL on an ELB.

Answer: D

QUESTION 270
A company is storing data on Amazon Simple Storage Service (S3). The company's security policy mandates that data is encrypted at rest. Which of the following methods can achieve this? Choose 3 answers

A.	Use Amazon S3 server-side encryption with AWS Key Management Service managed keys.
B.	Use Amazon S3 server-side encryption with customer-provided keys.

C.	Use Amazon S3 server-side encryption with EC2 key pair.
D.	Use Amazon S3 bucket policies to restrict access to the data at rest.
E.	Encrypt the data on the client-side before ingesting to Amazon S3 using their own master key.
F.	Use SSL to encrypt the data while in transit to Amazon S3.

Answer: ABE

QUESTION 271
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 idle. Once data is processed, change the storage class of the S3 objects to Reduced Redundancy Storage.
B.	Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SOS Once data is processed,
C.	Change the storage class of the S3 objects to Reduced Redundancy Storage. 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.
D.	Use SNS to pass job messages use Cloud Watch alarms to terminate spot worker instances when they become idle. Once data is processed, change the storage class of the S3 object to Glacier.

Answer: D

QUESTION 272
You control access to S3 buckets and objects with:

A.	Identity and Access Management (IAM) Policies.
B.	Access Control Lists (ACLs).
C.	Bucket Policies.
D.	All of the above

Answer: D

QUESTION 273
The AWS IT infrastructure that AWS provides, complies with the following IT security standards, including:

A.	SOC 1/SSAE 16/ISAE 3402 (formerly SAS 70 Type II), SOC 2 and SOC 3
B.	FISMA, DIACAP, and FedRAMP
C.	PCI DSS Level 1, ISO 27001, ITAR and FIPS 140-2
D.	HIPAA, Cloud Security Alliance (CSA) and Motion Picture Association of America (MPAA)
E.	All of the above

Answer: ABC

QUESTION 274
Auto Scaling requests are signed with a _____ signature calculated from the request and the user's private key.

A.	SSL
B.	AES-256
C.	HMAC-SHA1
D.	X.509

Answer: C

QUESTION 275
The following policy can be attached to an IAM group. It lets an IAM user in that group access a "home directory" in AWS S3 that matches their user name using the console.

```
{
"Version": "2012-10-17",
"Statement": [
{
"Action": ["s3:*"],
"Effect": "Allow",
"Resource": ["arn:aws:s3:::bucket-name"],
"Condition":{"StringLike":{"s3:prefix":["home/${aws:username}/*"]}} },
{
"Action":["s3:*"],
"Effect":"Allow",
"Resource": ["arn:aws:s3:::bucket-name/home/${aws:username}/*"] }
]
}
```

A.	True
B.	False

Answer: B

QUESTION 276
What does elasticity mean to AWS?

A.	The ability to scale computing resources up easily, with minimal friction and down with latency.
B.	The ability to scale computing resources up and down easily, with minimal friction.
C.	The ability to provision cloud computing resources in expectation of future demand.
D.	The ability to recover from business continuity events with minimal friction.

Answer: B

QUESTION 277
The following are AWS Storage services? Choose 2 Answers

A.	AWS Relational Database Service (AWS RDS)
B.	AWS ElastiCache
C.	AWS Glacier
D.	AWS Import/Export

Answer: BD

QUESTION 278
How is AWS readily distinguished from other vendors in the traditional IT computing landscape?

A.	Experienced. Scalable and elastic. Secure. Cost-effective. Reliable
B.	Secure. Flexible. Cost-effective. Scalable and elastic. Global
C.	Secure. Flexible. Cost-effective. Scalable and elastic. Experienced
D.	Flexible. Cost-effective. Dynamic. Secure. Experienced.

Answer: C

QUESTION 279
Your company is storing millions of sensitive transactions across thousands of 100-GB files that must be encrypted in transit and at rest. Analysts concurrently depend on subsets of files, which can consume up to 5 TB of space, to generate simulations that can be used to steer business decisions. You are required to design an AWS solution that can cost effectively accommodate the long-term storage and in-flight subsets of data.

A.	Use Amazon Simple Storage Service (S3) with server-side encryption, and run simulations on subsets in ephemeral drives on Amazon EC2.
B.	Use Amazon S3 with server-side encryption, and run simulations on subsets in-memory on Amazon EC2.
C.	Use HDFS on Amazon EMR, and run simulations on subsets in ephemeral drives on Amazon EC2.
D.	Use HDFS on Amazon Elastic MapReduce (EMR), and run simulations on subsets in-memory on Amazon Elastic Compute Cloud (EC2).
E.	Store the full data set in encrypted Amazon Elastic Block Store (EBS) volumes, and regularly capture snapshots that can be cloned to EC2 workstations.

Answer: D

QUESTION 280
Can you configure multiple Load Balancers with a single Auto Scaling group?

A.	No
B.	Yes, you can but only if it is configured with Amazon Redshift.
C.	Yes, you can provided the ELB is configured with Amazon AppStream.
D.	Yes

Answer: D
Explanation:
Yes, you can configure more than one load balancer with an autoscaling group. Auto Scaling integrates with Elastic Load Balancing to enable you to attach one or more load balancers to an

existing Auto Scaling group. After you attach the load balancer, it automatically registers the instances in the group and distributes incoming traffic across the instances.
http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html

QUESTION 281
An Auto Scaling group is running at the desired capacity of 5 instances and receives a trigger from the Cloudwatch Alarm to increase the capacity by 1. The cool down period is 5 minutes. Cloudwatch sends another trigger after 2 minutes to decrease the desired capacity by 1. What will be the count of instances at the end of 4 minutes?

A.	4
B.	5
C.	6
D.	7

Answer: C
Explanation:
The cool down period is the time difference between the end of one scaling activity (can be start or terminate) and the start of another one (can be start or terminate). During the cool down period, Auto Scaling does not allow the desired capacity of the Auto Scaling group to be changed by any other CloudWatch alarm. Thus, in this case the trigger from the second alarm will have no effect.
http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html#healthcheck

QUESTION 282
Which of the following is NOT a true statement about Auto Scaling?

A.	Auto Scaling can launch instances in different Azs.
B.	Auto Scaling can work with CloudWatch.
C.	Auto Scaling can launch an instance at a specific time.
D.	Auto Scaling can launch instances in different regions.

Answer: D
Explanation:
Auto Scaling provides an option to scale up and scale down based on certain conditions or triggers from Cloudwatch. A user can configure such that Auto Scaling launches instances across Azs, but it cannot span across regions.
http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-dg.pdf

QUESTION 283
A user wants to configure AutoScaling which scales up when the CPU utilization is above 70% and scales down when the CPU utilization is below 30%. How can the user configure AutoScaling for the above mentioned condition?

A.	Configure ELB to notify AutoScaling on load increase or decrease
B.	Use AutoScaling with a schedule
C.	Use AutoScaling by manually modifying the desired capacity during a condition
D.	Use dynamic AutoScaling with a policy

Answer: D
Explanation:

The user can configure the AutoScaling group to automatically scale up and then scale down based on the specified conditions. To configure this, the user must setup policies which will get triggered by the CloudWatch alarms.
http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-scale-based-on-demand.html

QUESTION 284
Can a user configure a custom health check with Auto Scaling?

A.	Yes, but the configured data will not be saved to Auto Scaling.
B.	No, only an ELB health check can be configured with Auto Scaling.
C.	Yes
D.	No

Answer: C
Explanation:
Auto Scaling can determine the health status of an instance using custom health checks. If you have custom health checks, you can send the information from your health checks to Auto Scaling so that Auto Scaling can use this information. For example, if you determine that an instance is not functioning as expected, you can set the health status of the instance to Unhealthy. The next time that Auto Scaling performs a health check on the instance, it will determine that the instance is unhealthy and then launch a replacement instance.
http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/healthcheck.html

QUESTION 285
You have setup an Auto Scaling group. The cool down period for the Auto Scaling group is 7 minutes. The first scaling activity request for the Auto Scaling group is to launch two instances. It receives the activity question at time "t", and the first instance is launched at t+3 minutes, while the second instance is launched at t+4 minutes. How many minutes after time "t" will Auto Scaling accept another scaling activity request?

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

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

QUESTION 286
You have just added a new instance to your Auto Scaling group, which receives ELB health checks. An ELB heath check says the new instance's state is OutOfService. What does Auto Scaling do in this particular scenario?

A.	It replaces the instance with a healthy one

B. It stops the instance
C. It marks an instance as unhealthy
D. It terminates the instance

Answer: C
Explanation:
If you have attached a load balancer to your Auto Scaling group, you can have Auto Scaling include the results of Elastic Load Balancing health checks when it determines the health status of an instance. After you add ELB health checks, Auto Scaling will mark an instance as unhealthy if Elastic Load Balancing reports the instance state as OutOfService. Frequently, an Auto Scaling instance that has just come into service needs to warm up before it can pass the Auto Scaling health check. Auto Scaling waits until the health check grace period ends before checking the health status of the instance. While the EC2 status checks and ELB health checks can complete before the health check grace period expires, Auto Scaling does not act on them until the health check grace period expires. To provide ample warm-up time for your instances, ensure that the health check grace period covers the expected startup time for your application.
http://docs.aws.amazon.com/autoscaling/latest/userguide/healthcheck.html

QUESTION 287
A sys admin is maintaining an application on AWS. The application is installed on EC2 and user has configured ELB and Auto Scaling. Considering future load increase, the user is planning to launch new servers proactively so that they get registered with ELB. How can the user add these instances with Auto Scaling?

A. Decrease the minimum limit of the Auto Scaling group
B. Increase the maximum limit of the Auto Scaling group
C. Launch an instance manually and register it with ELB on the fly
D. Increase the desired capacity of the Auto Scaling group

Answer: D
Explanation:
A user can increase the desired capacity of the Auto Scaling group and Auto Scaling will launch a new instance as per the new capacity. The newly launched instances will be registered with ELB if Auto Scaling group is configured with ELB. If the user decreases the minimum size the instances will be removed from Auto Scaling. Increasing the maximum size will not add instances but only set the maximum instance cap.
http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-manual-scaling.html

QUESTION 288
You have set up Auto Scaling to automatically scale in. Consequently, you must decide which instances Auto Scaling should end first. What should you use to configure this?

A. An Elastic Load Balancer
B. A termination policy
C. An IAM role
D. Another scaling group

Answer: B
Explanation:
If you instruct Auto Scaling to automatically scale in, you must decide which instances Auto Scaling should terminate first. This can be configured through the use of a termination policy.
http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AutoScalingBehavior.Insta

nceTermination.html

QUESTION 289
Which of the following commands accepts binary data as parameters?

A.	--user-data
B.	--ciphertext-key
C.	--aws-customer-key
D.	--describe-instances-user

Answer: A
Explanation:
For commands that take binary data as a parameter, specify that the data is binary content by using the fileb:// prefix.
Commands that accept binary data include:
aws ec2 run-instances --user-data parameter.
aws s3api put-object --sse-customer-key parameter.
aws kms decrypt --ciphertext-blob parameter.
http://docs.aws.amazon.com/cli/latest/userguide/aws-cli.pdf

QUESTION 290
To scale out the AWS resources using manual AutoScaling, which of the below mentioned parameters should the user change?

A.	Current capacity
B.	Desired capacity
C.	Preferred capacity
D.	Maximum capacity

Answer: B
Explanation:
The Manual Scaling as part of Auto Scaling allows the user to change the capacity of Auto Scaling group. The user can add / remove EC2 instances on the fly. To execute manual scaling, the user should modify the desired capacity. AutoScaling will adjust instances as per the requirements.
http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-manual-scaling.html

QUESTION 291
After moving an E-Commerce website for a client from a dedicated server to AWS you have also set up auto scaling to perform health checks on the instances in your group and replace instances that fail these checks. Your client has come to you with his own health check system that he wants you to use as it has proved to be very useful prior to his site running on AWS. What do you think would be an appropriate response to this given all that you know about auto scaling and CloudWatch?

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

information directly from your system to CloudWatch.

Answer: D
Explanation:
Auto Scaling periodically performs health checks on the instances in your group and replaces instances that fail these checks. By default, these health checks use the results of EC2 instance status checks to determine the health of an instance. If you use a load balancer with your Auto Scaling group, you can optionally choose to include the results of Elastic Load Balancing health checks.
Auto Scaling marks an instance unhealthy if the calls to the Amazon EC2 action DescribeInstanceStatus returns any other state other than running, the system status shows impaired, or the calls to Elastic Load Balancing action DescribeInstanceHealth returns OutOfService in the instance state field.
After an instance is marked unhealthy because of an Amazon EC2 or Elastic Load Balancing health check, it is scheduled for replacement.
You can customize the health check conducted by your Auto Scaling group by specifying additional checks or by having your own health check system and then sending the instance's health information directly from your system to Auto Scaling.
http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/healthcheck.html

QUESTION 292
Identify a benefit of using Auto Scaling for your application.

A. Your application gains better fault tolerance.
B. Your application optimizes only logistics and operations.
C. Your application receives latency requirements in every region.
D. You acquire clarity on prototypes in your application.

Answer: A
Explanation:
When you use Auto Scaling, your applications gain better fault tolerance. Auto Scaling can detect when an instance is unhealthy, terminate it, and launch an instance to replace it. You can also configure Auto Scaling to use multiple Availability Zones. If one Availability Zone becomes unavailable, Auto Scaling can launch instances in another one to compensate.
http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/how-as-works.html

QUESTION 293
A user has suspended the scaling process on the Auto Scaling group. A scaling activity to increase the instance count was already in progress. What effect will the suspension have on that activity?

A. No effect. The scaling activity continues
B. Pauses the instance launch and launches it only after Auto Scaling is resumed
C. Terminates the instance
D. Stops the instance temporary

Answer: A
Explanation:
The user may want to stop the automated scaling processes on the Auto Scaling groups either to perform manual operations or during emergency situations. To perform this, the user can suspend one or more scaling processes at any time. When this process is suspended, Auto Scaling creates no new scaling activities for that group. Scaling activities that were already in

progress before the group was suspended continue until completed.
http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html

QUESTION 294
Does Autoscaling automatically assign tags to resources?

A. No, not unless they are configured via API.
B. Yes, it does.
C. Yes, by default.
D. No, it does not.

Answer: B
Explanation:
Tags don't have any semantic meaning to Amazon EC2 and are interpreted strictly as a string of characters.
Tags are assigned automatically to the instances created by an Auto Scaling group. Auto Scaling adds a tag to the instance with a key of aws:autoscaling:groupName and a value of the name of the Auto Scaling group.
http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Using_Tags.html

QUESTION 295
If you have a running instance using an Amazon EBS boot partition, you can call the _____ API to release the compute resources but preserve the data on the boot partition.

A. Stop Instances
B. Terminate Instances
C. AMI Instance
D. Ping Instance

Answer: A
Explanation:
If you have a running instance using an Amazon EBS boot partition, you can also call the Stop Instances API to release the compute resources but preserve the data on the boot partition.
https://aws.amazon.com/ec2/faqs/#How_quickly_will_systems_be_running

QUESTION 296
Which EC2 functionality allows the user to place the Cluster Compute instances in clusters?

A. Cluster group
B. Cluster security group
C. GPU units
D. Cluster placement group

Answer: D
Explanation:
The Amazon EC2 cluster placement group functionality allows users to group cluster compute instances in clusters.
https://aws.amazon.com/ec2/faqs/

QUESTION 297

A user has launched a dedicated EBS backed instance with EC2. You are curious where the EBS volume for this instance will be created. Which statement is correct about the EBS volume's creation?

A.	The EBS volume will not be created on the same tenant hardware assigned to the dedicated instance
B.	AWS does not allow a dedicated EBS backed instance launch
C.	The EBS volume will be created on the same tenant hardware assigned to the dedicated instance
D.	The user can specify where the EBS will be created

Answer: A
Explanation:
The dedicated instances are Amazon EC2 instances that run in a Virtual Private Cloud (VPC) on hardware that is dedicated to a single customer. When a user launches an Amazon EBS-backed dedicated instance, the EBS volume does not run on single-tenant hardware.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/dedicated-instance.html

QUESTION 298
Which system is used by Amazon Machine Images paravirtual (PV) virtualization during the boot process?

A.	PV-BOOT
B.	PV-AMI
C.	PV-WORM
D.	PV-GRUB

Answer: D
Explanation:
Amazon Machine Images that use paravirtual (PV) virtualization use a system called PV-GRUB during the boot process. PV-GRUB is a paravirtual boot loader that runs a patched version of GNU GRUB 0.97. When you start an instance, PV-GRUB starts the boot process and then chain loads the kernel specified by your image's menu.lst file.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/UserProvidedKernels.html

QUESTION 299
You have written a CloudFormation template that creates 1 Elastic Load Balancer fronting 2 EC2 Instances. Which section of the template should you edit so that the DNS of the load balancer is returned upon creation of the stack?

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

Answer: B
Explanation:
You can use AWS CloudFormation's sample templates or create your own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run your application.
In the following example, the output named BackupLoadBalancerDNSName returns the DNS name for the resource with the logical ID BackupLoadBalancer only when the CreateProdResources condition is true. (The second output shows how to specify multiple

outputs.)
"Outputs" : {
"BackupLoadBalancerDNSName" : {
"Description": "The DNSName of the backup load balancer", "Value" : { "Fn::GetAtt" :
["BackupLoadBalancer", "DNSName"]}, "Condition" : "CreateProdResources"
},
"InstanceID" : {
"Description": "The Instance ID",
"Value" : { "Ref" : "EC2Instance" }
}
}
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/outputs-section-structure.htmlxx

QUESTION 300
In CloudFormation, if you want to map an Amazon Elastic Block Store to an Amazon EC2 instance, _____.

A.	you reference the logical IDs to associate the block stores with the instance
B.	you reference the physical IDs of the instance along with the resource type
C.	you reference the instance IDs of the block store along with the resource properties
D.	you reference the physical IDs of both the block stores and the instance

Answer: A
Explanation:
In AWS CloudFormation, if you want to map an Amazon Elastic Block Store to an Amazon EC2 instance, you reference the logical IDs to associate the block stores with the instance.
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/concept-resources.html

QUESTION 301
An organization hosts an app on EC2 instances which multiple developers need access to in order to perform updates.
The organization plans to implement some security best practices related to instance access.
Which one of the following recommendations will not help improve its security in this way?

A.	Disable the password based login for all the users. All the users should use their own keys to connect with the instance securely.
B.	Create an IAM policy allowing only IAM users to connect to the EC2 instances with their own SSH key.
C.	Create a procedure to revoke the access rights of the individual user when they are not required to connect to EC2 instance anymore for the purpose of application configuration.
D.	Apply the latest patch of OS and always keep it updated.

Answer: B
Explanation:
Since AWS is a public cloud any application hosted on EC2 is prone to hacker attacks. It becomes extremely important for a user to setup a proper security mechanism on the EC2 instances. A few of the security measures are listed below:
. Always keep the OS updated with the latest patch
. Always create separate users with in OS if they need to connect with the EC2 instances, create their keys and disable their password
. Create a procedure using which the admin can revoke the access of the user when the business work on the EC2 instance is completed

. Lock down unnecessary ports
. Audit any proprietary applications that the user may be running on the EC2 instance . Provide temporary escalated privileges, such as sudo for users who need to perform occasional privileged tasks
IAM is useful when users are required to work with AWS resources and actions, such as launching an instance. It is not useful in this case because it does not manage who can connect via RDP or SSH with an instance.
http://aws.amazon.com/articles/1233/

QUESTION 302
A user has configured two security groups which allow traffic as given below:
1: SecGrp1:
Inbound on port 80 for 0.0.0.0/0
Inbound on port 22 for 0.0.0.0/0
2: SecGrp2:
Inbound on port 22 for 10.10.10.1/32
If both the security groups are associated with the same instance, which of the below mentioned statements is true?

A.	It is not possible to have more than one security group assigned to a single instance
B.	It is not possible to create the security group with conflicting rules. AWS will reject the request
C.	It allows inbound traffic for everyone on both ports 22 and 80
D.	It allows inbound traffic on port 22 for IP 10.10.10.1 and for everyone else on port 80

Answer: C
Explanation:
A user can attach more than one security group to a single EC2 instance. In this case, the rules from each security group are effectively aggregated to create one set of rules. AWS uses this set of rules to determine whether to allow access or not. Thus, here the rule for port 22 with IP 10.10.10.1/32 will merge with IP 0.0.0.0/0 and open ports 22 and 80 for all.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html

QUESTION 303
You have a website which requires international presence and consequently you have set it up as follows.
It is hosted on 30 EC2 instances.
It is on in 15 regions around the globe.
Each region has 2 instances.
All the instances are a public hosted zone.
Which of the following is the best way to configure your site to maintain availability with minimum downtime if one of the 15 regions was to lose network connectivity for an extended period?
(Choose 2 answers)

A.	Create a Route 53 Latency Based Routing Record set that resolves to an Elastic Load Balancer in each region and has the Evaluate Target Health flag set to true.
B.	Create a Route 53 failover routing policy and configure an active-passive failover.
C.	Create a Route 53 Failover Routing Policy and assign each resource record set a unique identifier and a relative weight.
D.	Create a Route 53 Geolocation Routing Policy that resolves to an Elastic Load Balancer in each region and has the Evaluate Target Health flag set to false.

Answer: AB
Explanation:

It is best to use the latency routing policy when you have resources in multiple Amazon EC2 data centers that perform the same function and you want Amazon Route 53 to respond to DNS queries with the resources that provide the best latency. You could also use the failover routing policy (for public hosted zones only) when you want to configure an active-passive failover, in which one resource takes all traffic when it's available and the other resource takes all traffic when the first resource isn't available.
http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html#routing-policy-latency

QUESTION 304
A user is accessing an EC2 instance on the SSH port for IP 10.20.30.40/32. Which one is a secure way to configure that the instance can be accessed only from this IP?

A.	In the security group, open port 22 for IP 10.20.30.40
B.	In the security group, open port 22 for IP 10.20.30.0
C.	In the security group, open port 22 for IP 10.20.30.40/32
D.	In the security group, open port 22 for IP 10.20.30.40/0

Answer: C
Explanation:
In AWS EC2, while configuring a security group, the user needs to specify the IP address in CIDR notation. The CIDR IP range 10.20.30.40/32 says it is for a single IP 10.20.30.40. If the user specifies the IP as 10.20.30.40 only, the security group will not accept and ask for it in a CIDR format.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html

QUESTION 305
While assigning a tag to an instance, which of the below mentioned options is not a valid tag key/value pair?

A.	Key : "aws" Value:"aws"
B.	Key: "aws:name" Value: "instanceAnswer: Aws"
C.	Key: "Name :aws" Value: "instanceAnswer: Aws"
D.	Key : "nameAnswer: Aws" Value:"aws:instance"

Answer: B
Explanation:
In Amazon Web Services, to help manage EC2 instances as well their usage in a better way, the user can tag the instances. The tags are metadata assigned by the user which consists of a key and value. The tag key cannot have a prefix as "aws:", although it can have only "aws".
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html

QUESTION 306
Will you be able to access EC2 snapshots using the regular Amazon S3 APIs?

A.	Yes, you will be able to access using S3 APIs if you have chosen the snapshot to be stored in S3.
B.	No, snapshots are only available through the Amazon EBS APIs.
C.	Yes, you will be able to access them using S3 APIs as all snapshots are stored in S3.
D.	No, snapshots are only available through the Amazon EC2 APIs.

Answer: D

Explanation:
No, snapshots are only available through the Amazon EC2 APIs.
https://aws.amazon.com/ec2/faqs/

QUESTION 307
A user has created an AWS AMI. The user wants the AMI to be available only to his friend and not anyone else. How can the user manage this?

A.	Share the AMI with the community and setup the approval workflow before anyone launches it.
B.	It is not possible to share the AMI with the selected user.
C.	Share the AMI with a friend's AWS account ID.
D.	Share the AMI with a friend's AWS login ID.

Answer: C
Explanation:
In Amazon Web Services, If a user has created an AMI and wants to share with his friends and colleagues he can share the AMI with their AWS account ID. Once the AMI is shared the other user can access it from the community AMIs under private AMIs options.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/sharingamis-explicit.html

QUESTION 308
A user is planning to launch multiple EC2 instance same as current running instance. Which of the below mentioned parameters is not copied by Amazon EC2 in the launch wizard when the user has selected the option "Launch more like this"?

A.	Termination protection
B.	Tenancy setting
C.	Storage
D.	Shutdown behaviour

Answer: C
Explanation:
The Amazon EC2 console provides a "Launch more like this" wizard option that enables the user to use a current instance as a template for launching other instances. This option automatically populates the Amazon EC2 launch wizard with certain configuration details from the selected instance.
The following configuration details are copied from the selected instance into the launch wizard:
AMI ID
Instance type
Availability Zone, or the VPC and subnet in which the selected instance is located Public IPv4 address. If the selected instance currently has a public IPv4 address, the new instance receives a public IPv4 address - regardless of the selected instance's default public IPv4 address setting. For more information about public IPv4 addresses, see Public IPv4 Addresses and External DNS Hostnames.
Placement group, if applicable
IAM role associated with the instance, if applicable Shutdown behavior setting (stop or terminate)
Termination protection setting (true or false)
CloudWatch monitoring (enabled or disabled)
Amazon EBS-optimization setting (true or false)
Tenancy setting, if launching into a VPC (shared or dedicated) Kernel ID and RAM disk ID, if applicable
User data, if specified
Tags associated with the instance, if applicable

Security groups associated with the instance
The following configuration details are not copied from your selected instance; instead, the wizard applies their default settings or behavior:
(VPC only) Number of network interfaces: The default is one network interface, which is the primary network interface (eth0).
Storage: The default storage configuration is determined by the AMI and the instance type.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/launching-instance.html

QUESTION 309
A user has launched an EBS optimized instance with EC2. Which of the below mentioned options is the correct statement?

A.	It provides additional dedicated capacity for EBS IO
B.	The attached EBS will have greater storage capacity
C.	The user will have a PIOPS based EBS volume
D.	It will be launched on dedicated hardware in VPC

Answer: A
Explanation:
An Amazon EBS-optimized instance uses an optimized configuration stack and provides additional, dedicated capacity for the Amazon EBS I/O. This optimization provides the best performance for the user's Amazon EBS volumes by minimizing contention between the Amazon EBS I/O and other traffic from the user's instance.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSOptimized.html

QUESTION 310
Which status represents a failure state in AWS CloudFormation?

A.	ROLLBACK_IN_PROGRESS
B.	DELETE_IN_PROGRESS
C.	UPDATE_COMPLETE_CLEANUP_IN_PROGRESS
D.	REVIEW_IN_PROGRESS

Answer: A
Explanation:
ROLLBACK_IN_PROGRESS means an ongoing removal of one or more stacks after a failed stack creation or after an explicitly canceled stack creation. DELETE_IN_PROGRESS means an ongoing removal of one or more stacks. REVIEW_IN_PROGRESS means an ongoing creation of one or more stacks with an expected StackId but without any templates or resources. UPDATE_COMPLETE_CLEANUP_IN_PROGRESS means an ongoing removal of old resources for one or more stacks after a successful stack update.
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-describing-stacks.html

QUESTION 311
You are playing around with setting up stacks using JSON templates in CloudFormation to try and understand them a little better. You have set up about 5 or 6 but now start to wonder if you are being charged for these stacks. What is AWS's billing policy regarding stack resources?

A.	You are not charged for the stack resources if they are not taking any traffic.
B.	You are charged for the stack resources for the time they were operating (but not if you deleted the

stack within 30 minutes)

C. You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 60 minutes)

D. You are charged for the stack resources for the time they were operating (even if you deleted the stack right away)

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

QUESTION 312
In an AWS CloudFormation template, each resource declaration includes:

A. a logical ID, a resource type, and resource properties
B. a variable resource name and resource attributes
C. an IP address and resource entities
D. a physical ID, a resource file, and resource data

Answer: A
Explanation:
In AWS CloudFormation, each resource declaration includes three parts: a logical ID that is unique within the template, a resource type, and resource properties.
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/concept-resources.html

QUESTION 313
For AWS CloudFormation, which stack state refuses UpdateStack calls?

A. UPDATE_ROLLBACK_FAILED
B. UPDATE_ROLLBACK_COMPLETE
C. UPDATE_COMPLETE
D. CREATE_COMPLETE

Answer: A
Explanation:
When a stack is in the UPDATE_ROLLBACK_FAILED state, you can continue rolling it back to return it to a working state (to UPDATE_ROLLBACK_COMPLETE). You cannot update a stack that is in the UPDATE_ROLLBACK_FAILED state. However, if you can continue to roll it back, you can return the stack to its original settings and try to update it again.
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating-sta cks-continueupdaterollback.html

QUESTION 314
In the context of AWS CloudFormation, which of the following statements is correct?

A. Actual resource names are a combination of the resource ID, stack, and logical resource name.
B. Actual resource name is the stack resource name.
C. Actual resource name is the logical resource name.
D. Actual resource names are a combination of the stack and logical resource name.

Answer: D
Explanation:
In AWS CloudFormation, actual resource names are a combination of the stack and logical resource name. This allows multiple stacks to be created from a template without fear of name collisions between AWS resources.
https://aws.amazon.com/cloudformation/faqs/

QUESTION 315
When using the AWS CLI for AWS CloudFormation, which of the following commands returns a description of the specified resource in the specified stack?

A. describe-stack-events
B. describe-stack-resource
C. create-stack-resource
D. describe-stack-returns

Answer: B
Explanation:
aws-cli cloudformation
describe-stack-resource
Description
Returns a description of the specified resource in the specified stack. For deleted stacks, describe-stack-resource returns resource information for up to 90 days after the stack has been deleted.
http://docs.aws.amazon.com/cli/latest/reference/cloudformation/describe-stack-resource.html

QUESTION 316
A user is using CloudFormation to launch an EC2 instance and then configure an application after the instance is launched. The user wants the stack creation of ELB and AutoScaling to wait until the EC2 instance is launched and configured properly. How can the user configure this?

A. The user can use the DependentCondition resource to hold the creation of the other dependent resources.
B. It is not possible that the stack creation will wait until one service is created and launched.
C. The user can use the HoldCondition resource to wait for the creation of the other dependent resources.
D. The user can use the WaitCondition resource to hold the creation of the other dependent resources.

Answer: D
Explanation:
AWS CloudFormation is an application management tool that provides application modeling, deployment, configuration, management, and related activities. AWS CloudFormation provides a WaitCondition resource that acts as a barrier and blocks the creation of other resources until a completion signal is received from an external source, such as a user application or management system.
http://aws.amazon.com/cloudformation/faqs/

QUESTION 317
AWS _____ supports _____ environments as one of the AWS resource types.

A.	Elastic Beanstalk; Elastic Beanstalk application
B.	CloudFormation; Elastic Beanstalk application
C.	Elastic Beanstalk ; CloudFormation application
D.	CloudFormation; CloudFormation application

Answer: B
Explanation:
AWS CloudFormation and AWS Elastic Beanstalk services are designed to complement each other. AWS CloudFormation supports Elastic Beanstalk application environments as one of the AWS resource types.
http://aws.amazon.com/cloudformation/faqs/

QUESTION 318
AWS CloudFormation _____ are special actions you use in your template to assign values to properties that are not available until runtime.

A.	intrinsic functions
B.	properties declarations
C.	output functions
D.	conditions declarations

Answer: A
Explanation:
AWS CloudFormation intrinsic functions are special actions you use in your template to assign values to properties not available until runtime. Each function is declared with a name enclosed in quotation marks (""), a single colon, and its parameters.
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-fuctions-structure.html

QUESTION 319
For Amazon EC2 issues, while troubleshooting AWS CloudFormation, you need to view the cloud-init and cfn logs for more information. Identify a directory to which these logs are published.

A.	/var/opt/log/ec2
B.	/var/log/lastlog
C.	/var/log/
D.	/var/log/ec2

Answer: C
Explanation:
When you use AWS CloudFormation, you might encounter issues when you create, update, or delete AWS CloudFormation stacks.
For Amazon EC2 issues, view the cloud-init and cfn logs. These logs are published on the Amazon EC2 instance in the /var/log/ directory. These logs capture processes and command outputs while AWS CloudFormation is setting up your instance. For Windows, view the EC2Configure service and cfn logs in %ProgramFiles%\Amazon\EC2ConfigService and C:\cfn\log.

You can also configure your AWS CloudFormation template so that the logs are published to Amazon CloudWatch, which displays logs in the AWS Management Console so you don't have to connect to your Amazon EC2 instance.
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/troubleshooting.html

QUESTION 320
True or false: In a CloudFormation template, you can reuse the same logical ID several times to reference the resources in other parts of the template.

A.	True, a logical ID can be used several times to reference the resources in other parts of the template.
B.	False, a logical ID must be unique within the template.
C.	False, you can mention a resource only once and you cannot reference it in other parts of a template.
D.	False, you cannot reference other parts of the template.

Answer: B
Explanation:
In AWS CloudFormation, the logical ID must be alphanumeric (A-Za-z0-9) and unique within the template. You use the logical name to reference the resource in other parts of the template.
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/concept-resources.html

QUESTION 321
True or false: In CloudFormation, you cannot create an Amazon RDS DB instance from a snapshot.

A.	False, you can specify it in attributes
B.	False, you can specify it in condition
C.	False, you can specify it in resource properties
D.	True

Answer: C
Explanation:
In AWS CloudFormation, resource properties are additional options that you can specify on a resource. For example, you can specify the DB snapshot property for an Amazon RDS DB instance in order to create a DB instance from a snapshot.
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/concept-resources.html

QUESTION 322
How can you check the operational validity of your AWS CloudFormation template?

A.	To check the operational validity, you need to attempt to create the stack.
B.	There is no way to check the operational validity of your AWS CloudFormation template.
C.	To check the operational validity, you need a sandbox or test area for AWS CloudFormation stacks.
D.	To check the operational validity, you need to use the aws cloudformation validate-template command.

Answer: A
Explanation:
In AWS CloudFormation, to check the operational validity, you need to attempt to create the

stack. There is no sandbox or test area for AWS CloudFormation stacks, so you are charged for the resources you create during testing.
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-validate-template.html

QUESTION 323
What is a circular dependency in AWS CloudFormation?

A.	When Nested Stacks depend on each other.
B.	When Resources form a DependOn loop.
C.	When a Template references an earlier version of itself.
D.	When a Template references a region, which references the original Template.

Answer: B
Explanation:
To resolve a dependency error, add a DependsOn attribute to resources that depend on other resources in your template. In some cases, you must explicitly declare dependencies so that AWS CloudFormation can create or delete resources in the correct order. For example, if you create an Elastic IP and a VPC with an Internet gateway in the same stack, the Elastic IP must depend on the Internet gateway attachment. For additional information, see DependsOn Attribute.
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/troubleshooting.html#troubleshooting-errors-dependency-error

QUESTION 324
You need to develop and run some new applications on AWS and you know that Elastic Beanstalk and CloudFormation can both help as a deployment mechanism for a broad range of AWS resources. Which of the following is TRUE statements when describing the differences between Elastic Beanstalk and CloudFormation?

A.	AWS Elastic Beanstalk introduces two concepts: The template, a JSON or YAML-format, text-based file
B.	Elastic Beanstalk supports AWS CloudFormation application environments as one of the AWS resource types.
C.	Elastic Beanstalk automates and simplifies the task of repeatedly and predictably creating groups of related resources that power your applications. CloudFormation does not.
D.	You can design and script custom resources in CloudFormation

Answer: D
Explanation:
These services are designed to complement each other. AWS Elastic Beanstalk provides an environment to easily deploy and run applications in the cloud. It is integrated with developer tools and provides a one-stop experience for you to manage the lifecycle of your applications. AWS CloudFormation is a convenient provisioning mechanism for a broad range of AWS resources. It supports the infrastructure needs of many different types of applications such as existing enterprise applications, legacy applications, applications built using a variety of AWS resources and container-based solutions (including those built using AWS Elastic Beanstalk). AWS CloudFormation supports Elastic Beanstalk application environments as one of the AWS resource types. This allows you, for example, to create and manage an AWS Elastic Beanstalk-hosted application along with an RDS database to store the application data. In addition to RDS instances, any other supported AWS resource can be added to the group as well.
https://aws.amazon.com/cloudformation/faqs/

QUESTION 325
An elastic network interface (ENI) is a virtual network interface that you can attach to an instance in a VPC. An ENI can include one public IP address, which can be auto-assigned to the elastic network interface for eth0 when you launch an instance, but only when you _____.

A.	create an elastic network interface for eth1
B.	include a MAC address
C.	use an existing network interface
D.	create an elastic network interface for eth0

Answer: D
Explanation:
An elastic network interface (ENI) is defined as a virtual network interface that you can attach to an instance in a VPC and can include one public IP address, which can be auto-assigned to the elastic network interface for eth0 when you launch an instance, but only when you create an elastic network interface for eth0 instead of using an existing network interface.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 326
After setting an AWS Direct Connect, which of the following cannot be done with an AWS Direct Connect Virtual Interface?

A.	You can exchange traffic between the two ports in the same region connecting to different Virtual Private Gateways (VGWs) if you have more than one virtual interface.
B.	You can change the region of your virtual interface.
C.	You can delete a virtual interface; if its connection has no other virtual interfaces, you can delete the connection.
D.	You can create a hosted virtual interface.

Answer: A
Explanation:
You must create a virtual interface to begin using your AWS Direct Connect connection. You can create a public virtual interface to connect to public resources or a private virtual interface to connect to your VPC. Also, it is possible to configure multiple virtual interfaces on a single AWS Direct Connect connection, and you'll need one private virtual interface for each VPC to connect to. Each virtual interface needs a VLAN ID, interface IP address, ASN, and BGP key. To use your AWS Direct Connect connection with another AWS account, you can create a hosted virtual interface for that account. These hosted virtual interfaces work the same as standard virtual interfaces and can connect to public resources or a VPC.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/WorkingWithVirtualInterfaces.html

QUESTION 327
Identify a correct statement about the expiration date of the "Letter of Authorization and Connecting Facility Assignment (LOA-CFA)," which lets you complete the Cross Connect step of setting up your AWS Direct Connect.

A.	If the cross connect is not completed within 90 days, the authority granted by the LOA-CFA expires.
B.	If the virtual interface is not created within 72 days, the LOA-CFA becomes outdated.
C.	If the cross connect is not completed within a user-defined time, the authority granted by the LOA-CFA expires.
D.	If the cross connect is not completed within the specified duration from the appropriate provider,

	the LOA-CFA expires.

Answer: A
Explanation:
An AWS Direct Connect location provides access to AWS in the region it is associated with. You can establish connections with AWS Direct Connect locations in multiple regions, but a connection in one region does not provide connectivity to other regions. Note: If the cross connect is not completed within 90 days, the authority granted by the LOA-CFA expires.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/Colocation.html

QUESTION 328
Which of the following is the final step that should be completed to start using AWS Direct Connect?

A.	Creating your Virtual Interface
B.	Configuring your router
C.	Completing the Cross Connect
D.	Verifying your Virtual Interface

Answer: D
Explanation:
You can get started using AWS Direct Connect by completing the following steps. Step 1: Sign Up for Amazon Web Services Step 2: Submit AWS Direct Connect Connection Request Step 3: Complete the Cross Connect (optional) Step 4: Configure Redundant Connections with AWS Direct Connect Step 5: Create a Virtual Interface Step 6: Download Router Configuration Step 7: Verify Your Virtual Interface
http://docs.aws.amazon.com/directconnect/latest/UserGuide/getstarted.html#connected

QUESTION 329
A user has created a VPC with CIDR 20.0.0.0/16. The user has created one subnet with CIDR 20.0.0.0/16 by mistake. The user is trying to create another subnet of CIDR 20.0.1.0/24. How can the user create the second subnet?

A.	The user can modify the first subnet CIDR with AWS CLI
B.	The user can modify the first subnet CIDR from the console
C.	There is no need to update the subnet as VPC automatically adjusts the CIDR of the first subnet based on the second subnet's CIDR
D.	It is not possible to create a second subnet with overlapping IP CIDR without deleting the first subnet.

Answer: D
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside the subnet. The user can create a subnet with the same size of VPC. However, he cannot create any other subnet since the CIDR of the second subnet will conflict with the first subnet. The user cannot modify the CIDR of a subnet once it is created. Thus, in this case if required, the user has to delete the subnet and create new subnets.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

QUESTION 330

Which of the following should be followed before connecting to Amazon Virtual Private Cloud (Amazon VPC) using AWS Direct Connect?

A.	Provide a public Autonomous System Number (ASN) to identify your network on the Internet.
B.	Create a virtual private gateway and attach it to your Virtual Private Cloud (VPC).
C.	Allocate a private IP address to your network in the 122.x.x.x range.
D.	Provide a public IP address for each Border Gateway Protocol (BGP) session.

Answer: B
Explanation:
To connect to Amazon Virtual Private Cloud (Amazon VPC) by using AWS Direct Connect, you must first do the following:
Provide a private Autonomous System Number (ASN) to identify your network on the Internet. Amazon then allocates a private IP address in the 169.x.x.x range to you. Create a virtual private gateway and attach it to your VPC.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

QUESTION 331
Your supervisor has given you the task of creating an elastic network interface on each of your web servers that connect to a mid-tier network where an application server resides. He also wants this set up as a Dual-homed Instance on Distinct Subnets. Instead of routing network packets through the dual-homed instances, where should each dual-homed instance receive and process requests to fulfil his criteria?

A.	On one of the web servers
B.	On the front end
C.	On the back end
D.	Through a security group

Answer: B
Explanation:
You can place an elastic network interface on each of your web servers that connects to a mid-tier network where an application server resides. The application server can also be dual-homed to a back-end network (subnet) where the database server resides. If it is set up like this, instead of routing network packets through the dual-homed instances, each dual-homed instance receives and processes requests on the front end and initiates a connection to the back end before finally sending requests to the servers on the back-end network.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 332
A user has created a VPC with CIDR 20.0.0.0/16 using the wizard. The user has created a public subnet CIDR (20.0.0.0/24) and VPN only subnets CIDR (20.0.1.0/24) along with the VPN gateway (vgw-123456) to connect to the user's data centre. The user's data centre has CIDR 172.28.0.0/12. The user has also setup a NAT instance (i-123456) to allow traffic to the internet from the VPN subnet. Which of the below mentioned options is not a valid entry for the main route table in this scenario?

A.	Destination: 20.0.0.0/16 and Target: local
B.	Destination: 0.0.0.0/0 and Target: i-123456
C.	Destination: 172.28.0.0/12 and Target: vgw-123456
D.	Destination: 20.0.1.0/24 and Target: i-123456

Answer: D
Explanation:
The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. When the user has configured this setup with Wizard, it will create a virtual private gateway to route all traffic of the VPN subnet. If the user has setup a NAT instance to route all the internet requests then all requests to the internet should be routed to it. All requests to the organization's DC will be routed to the VPN gateway. Here are the valid entries for the main route table in this scenario:
Destination: 0.0.0.0/0 & Target: i-123456 (To route all internet traffic to the NAT Instance)
Destination: 172.28.0.0/12 & Target: vgw-123456 (To route all the organization's data centre traffic to the VPN gateway)
Destination: 20.0.0.0/16 & Target: local (To allow local routing in VPC)
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario3.html

QUESTION 333
While implementing the policy keys in AWS Direct Connect, if you use _____ key and the request comes from an Amazon EC2 instance, the instance's public IP address is evaluated to determine if access is allowed.

A.	aws:SourceIp
B.	aws:EpochTime
C.	aws:CurrentTime
D.	aws:SecureTransport

Answer: A
Explanation:
If you use aws:SourceIp, and the request comes from an Amazon EC2 instance, the instance's public IP address is used to determine if access is allowed.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/using_iam.html

QUESTION 334
In which step of "start using AWS Direct Connect" steps is the virtual interface you created tagged with a customer-provided tag that complies with the Ethernet 802.1Q standard?

A.	Download Router Configuration.
B.	Complete the Cross Connect.
C.	Configure Redundant Connections with AWS Direct Connect.
D.	Create a Virtual Interface.

Answer: D
Explanation:
In the list of using Direct Connect steps, the create a Virtual Interface step is to provision your virtual interfaces. Each virtual interface must be tagged with a customer-provided tag that complies with the Ethernet 802.1Q standard. This tag is required for any traffic traversing the AWS Direct Connect connection.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/getstarted.html#createvirtualinterface

QUESTION 335
A user has created a VPC with CIDR 20.0.0.0/16 using the VPC wizard. The user has created public and VPN only subnets along with hardware VPN access to connect to the user's data

centre. The user has not yet launched any instance as well as modified or deleted any setup. He wants to delete this VPC from the console. Will the console allow the user to delete the VPC?

A.	Yes, the user can detach the virtual private gateway and then use the VPC console to delete the VPC.
B.	No, since the NAT instance is running, the user cannot delete the VPC.
C.	Yes, the user can use the CLI to delete the VPC that will detach the virtual private gateway automatically.
D.	No, the VPC console needs to be accessed using an administrator account to delete the VPC.

Answer: A
Explanation:
You can delete your VPC at any time (for example, if you decide it's too small). However, you must terminate all instances in the VPC first. When you delete a VPC using the VPC console, Amazon deletes all its components, such as subnets, security groups, network ACLs, route tables, Internet gateways, VPC peering connections, and DHCP options.If you have a VPN connection, you don't have to delete it or the other components related to the VPN (such as the customer gateway and virtual private gateway).
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html#VPC_Deleting

QUESTION 336
You have been asked to set up a public website on AWS with the following criteria:
You want the database and the application server running on an Amazon VPC. You want the database to be able to connect to the Internet so that it can be automatically updated to the correct patch level.
You do not want to receive any incoming traffic from the Internet to the database. Which solutions would be the best to satisfy all the above requirements for your planned public website on AWS? (Choose 2 answers)

A.	Set up both the public website and the database on a public subnet and block all incoming requests from the Internet with a Network Access Control List (NACL)
B.	Set up both the public website and the database on a public subnet, and block all incoming requests from the Internet with a security group which only allows access from the IP of the public website.
C.	Set up the public website on a public subnet and set up the database in a private subnet which connects to the Internet via a NAT instance.
D.	Set up both the public website and the database on a private subnet and block all incoming requests from the Internet with a Network Access Control List (NACL). Set up a Security group between the public website and the database which only allows access via port 80.

Answer: BC
Explanation:
You want the database to be able to connect to the Internet you need to either set it up on a public subnet or set it up on a private subnet which connects to the Internet via a NAT instance
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

QUESTION 337
Which statement is NOT true about accessing remote AWS region in the US by your AWS Direct Connect which is located in the US?

A.	AWS Direct Connect locations in the United States can access public resources in any US region.
B.	You can use a single AWS Direct Connect connection to build multi-region services.

C. Any data transfer out of a remote region is billed at the location of your AWS Direct Connect data transfer rate.

D. To connect to a VPC in a remote region, you can use a virtual private network (VPN) connection over your public virtual interface.

Answer: C
Explanation:
AWS Direct Connect locations in the United States can access public resources in any US region. You can use a single AWS Direct Connect connection to build multi-region services. To connect to a VPC in a remote region, you can use a virtual private network (VPN) connection over your public virtual interface.
To access public resources in a remote region, you must set up a public virtual interface and establish a border gateway protocol (BGP) session. Then your router learns the routes of the other AWS regions in the US. You can then also establish a VPN connection to your VPC in the remote region.
Any data transfer out of a remote region is billed at the remote region data transfer rate.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/remote_regions.html

QUESTION 338
Which of the following statements is NOT correct when working with your AWS Direct Connect connection after it is set up completely?

A. You can manage your AWS Direct Connect connections and view the connection details.

B. You can delete a connection as long as there are no virtual interfaces attached to it.

C. You cannot view the current connection ID and verify if it matches the connection ID on the Letter of Authorization (LOA).

D. You can accept a host connection by purchasing a hosted connection from the partner (APN).

Answer: C
Explanation:
You can manage your AWS Direct Connect connections and view connection details, accept hosted connections, and delete connections. You can view the current status of your connection. You can also view your connection ID, which looks similar to this example dxcon-xxxx, and verify that it matches the connection ID on the Letter of Authorization (LOA) that you received from Amazon.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/viewdetails.html

QUESTION 339
Over which of the following Ethernet standards does AWS Direct Connect link your internal network to an AWS Direct Connect location?

A. Single mode fiber-optic cable

B. Multi mode fiber-optic cable

C. Shielded balanced copper cable

D. Twisted pair cable

Answer: A
Explanation:
AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 1 gigabit or 10 gigabit Ethernet single mode fiber-optic cable.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

QUESTION 340
Your supervisor has given you the task of creating an elastic network interface on each of your web servers that connect to a mid-tier network where an application server resides. He also wants this set up as a Dual-homed Instance on Distinct Subnets. Instead of routing network packets through the dual-homed instances, where should each dual-homed instance receive and process requests to fulfil his criteria?

A.	On the front end
B.	On one of the web servers
C.	Through a security group
D.	On the back end

Answer: A
Explanation:
You can place an elastic network interface on each of your web servers that connects to a mid-tier network where an application server resides. The application server can also be dual-homed to a back-end network (subnet) where the database server resides. If it is set up like this, instead of routing network packets through the dual-homed instances, each dual-homed instance receives and processes requests on the front end and initiates a connection to the back end before finally sending requests to the servers on the back-end network.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 341
One of the components that is part of ec2-net-utils used with ENI's is ec2ifscan. Which of the following is not correct about ec2-net-utils?

A.	ec2-net-utils generates an interface configuration file suitable for use with DHCP.
B.	ec2-net-utils extends the functionality of the standard ifup.
C.	ec2-net-utils detaches a primary network interface from an instance.
D.	ec2-net-utils identifies network interfaces when they are attached, detached, or reattached to a running instance.

Answer: C
Explanation:
Each instance in a VPC has a default elastic network interface (the primary network interface) 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. You can create and attach additional elastic network interfaces. Amazon Linux AMIs may contain additional scripts installed by AWS, known as ec2-net-utils. One of the components that is part of ec2-net-utils used with ENI's is ec2ifscan. Its function is to check for network interfaces that have not been configured and configure them.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 342
A user wants to create a public subnet in VPC and launch an EC2 instance within it. The user has not selected the option to assign a public IP address while launching the instance. Which of the below mentioned statements is true with respect to this scenario?

A.	The instance will always have a public DNS attached to the instance by default
B.	The user would need to create a default route to IGW in subnet's route table and then attach an elastic IP to the instance to connect from the internet
C.	The user can directly attach an elastic IP to the instance

D. The instance will never launch if the public IP is not assigned

Answer: B
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. When the user is launching an instance he needs to select an option which attaches a public IP to the instance. If the user has not selected the option to attach the public IP then it will only have a private IP when launched. The user cannot connect to the instance from the internet. If the user wants an elastic IP to connect to the instance from the internet he should create an internet gateway and assign an elastic IP to instance.
http://docs.aws.amazon.com/AmazonVPC/latest/GettingStartedGuide/LaunchInstance.html

QUESTION 343
A user has created a VPC with a public subnet. The user has terminated all the instances which are part of the subnet. Which of the below mentioned statements is true with respect to this scenario?

A. The subnet to which the instances were launched with will be deleted
B. When the user launches a new instance it cannot use the same subnet
C. The user cannot delete the VPC since the subnet is not deleted
D. Secondary network interfaces attached to the terminated instances may persist.

Answer: D
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. When an instance is launched it will have a network interface attached with it. The user cannot delete the subnet until he terminates the instance and deletes the network interface. By default, network interfaces that are automatically created and attached to instances using the console are set to terminate when the instance terminates. However, network interfaces created using the command line interface aren't set to terminate when the instance terminates.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 344
When configuring your customer gateway to connect to your VPC, the _____ Association is established first between the virtual private gateway and customer gateway using the Pre-Shared Key as the authenticator.

A. Ipsec
B. BGP
C. IKE Security
D. Tunnel

Answer: C
Explanation:
When configuring your customer gateway to connect to your VPC, several steps need to be completed. The IKE Security Association is established first between the virtual private gateway and customer gateway using the Pre-Shared Key as the authenticator.
http://docs.aws.amazon.com/AmazonVPC/latest/NetworkAdminGuide/Introduction.html

QUESTION 345
An organization is trying to setup a VPC with Auto Scaling. Which configuration steps below is not required to setup AWS VPC with Auto Scaling?

A.	Configure the Auto Scaling group with the VPC ID in which instances will be launched.
B.	Configure the Auto Scaling Launch configuration with multiple subnets of the VPC to enable the Multi AZ feature.
C.	Configure the Auto Scaling Launch configuration which does not allow assigning a public IP to instances.
D.	Configure the Auto Scaling Launch configuration with the VPC security group.

Answer: B
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an Auto Scaling group. Before creating the Auto Scaling group it is recommended that the user creates the Launch configuration. Since it is a VPC, it is recommended to select the parameter which does not allow assigning a public IP to the instances.
The user should also set the VPC security group with the Launch configuration and select the subnets where the instances will be launched in the AutoScaling group. The HA will be provided as the subnets may be a part of separate AZs.
http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/autoscalingsubnets.html

QUESTION 346
An organization is planning to host a Wordpress blog as well as joomla CMS on a single instance launched with VPC. The organization wants to create separate domains for each application using Route 53. The organization may have about ten instances each with these two applications. While launching each instance, the organization configured two separate network interfaces (primary + secondary ENI) with their own Elastic IPs to the instance. The suggestion was to use a public IP from AWS instead of an Elastic IP as the number of elastic IPs allocation per region is restricted in the account. What action will you recommend to the organization?

A.	Only Elastic IP can be used by requesting limit increase, since AWS does not assign a public IP to an instance with multiple ENIs.
B.	AWS VPC does not attach a public IP to an ENI; so the only way is to use an Elastic IP.
C.	I agree with the suggestion but will prefer that the organization should use separate subnets with each ENI for different public IPs.
D.	I agree with the suggestion and it is recommended to use a public IP from AWS since the organization is going to use DNS with Route 53.

Answer: A
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. An Elastic Network Interface (ENI) is a virtual network interface that the user can attach to an instance in a VPC.
The user can attach upto two ENIs with a single instance. However, AWS cannot assign a public IP when there are two ENIs attached to a single instance. It is recommended to assign an elastic IP in this scenario. If the organization wants more than 5 EIPs they can request AWS to increase the number.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 347
A user has created a VPC with public and private subnets. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.1.0/24 and the public subnet uses CIDR 20.0.0.0/24. The user is planning to host a web server in the public subnet (port 80) and a DB server in the private subnet (port 3306). The user is configuring a security group of the NAT instance. Which of the below mentioned entries is not required in NAT's security group for the database servers to connect to the Internet for software updates ?

A.	For Outbound allow Destination: 0.0.0.0/0 on port 443
B.	For Inbound allow Source: 20.0.1.0/24 on port 80
C.	For Inbound allow Source: 20.0.0.0/24 on port 80
D.	For Outbound allow Destination: 0.0.0.0/0 on port 80

Answer: C
Explanation:
A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet to host the web server and DB server respectively, the user should configure that the instances in the private subnet can connect to the internet using the NAT instances. The user should first configure that NAT can receive traffic on ports 80 and 443 from the private subnet. Thus, allow ports 80 and 443 in Inbound for the private subnet 20.0.1.0/24. Now to route this traffic to the internet configure ports 80 and 443 in Outbound with destination 0.0.0.0/0. The NAT should not have an entry for the public subnet CIDR.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

QUESTION 348
A user has created a VPC with public and private subnets using the VPC wizard. Which of the below mentioned statements is true in this scenario?

A.	The user has to manually create a NAT instance
B.	The Amazon VPC will automatically create a NAT instance with the micro size only
C.	VPC updates the main route table used with the private subnet, and creates a custom route table with a public subnet
D.	VPC updates the main route table used with a public subnet, and creates a custom route table with a private subnet

Answer: C
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public subnet, the instances in the public subnet can receive inbound traffic directly from the internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create a NAT instance of a smaller or higher size, respectively. The VPC has an implied router and the VPC wizard updates the main route table used with the private subnet, creates a custom route table and associates it with the public subnet.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

QUESTION 349
A user has created a VPC with two subnets: one public and one private. The user is planning to run the patch update for the instances in the private subnet. How can the instances in the private subnet connect to the internet?

A. The private subnet can never connect to the internet
B. Use NAT with an elastic IP
C. Use the internet gateway with a private IP
D. Allow outbound traffic in the security group for port 80 to allow internet updates

Answer: B
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. If the user has created two subnets (one private and one public), they would need a Network Address Translation (NAT) instance with the elastic IP address. This enables the instances in the private subnet to send requests to the internet (for example, to perform software updates).
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

QUESTION 350
A user has created a VPC with public and private subnets using the VPC Wizard. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.0.0/24. Which of the below mentioned entries are required in the main route table to allow the instances in VPC to communicate with each other?

A. Destination : 20.0.0.0/0 and Target : ALL
B. Destination : 20.0.0.0/16 and Target : Local
C. Destination : 20.0.0.0/24 and Target : Local
D. Destination : 20.0.0.0/16 and Target : ALL

Answer: B
Explanation:
A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the Internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create two route tables and attach to the subnets. The main route table will have the entry "Destination: 20.0.0.0/16 and Target: Local", which allows all instances in the VPC to communicate with each other.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

QUESTION 351
You want to establish redundant VPN connections and customer gateways on your network by setting up a second VPN connection. Which of the following will ensure that this functions correctly?

A. The customer gateway IP address for the second VPN connection must be publicly accessible.
B. The virtual gateway IP address for the second VPN connection must be publicly accessible.
C. The customer gateway IP address for the second VPN connection must use dynamic routes.
D. The customer gateway IP address for the second VPN connection must be privately accessible and be the same public IP address that you are using for the first VPN connection.

Answer: A
Explanation:
To establish redundant VPN connections and customer gateways on your network, you would need to set up a second VPN connection. However, you must ensure that the customer gateway IP address for the second VPN connection is publicly accessible.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html

QUESTION 352
Someone is creating a VPC for their application hosting. He has created two private subnets in the same availability zone and created one subnet in a separate availability zone. He wants to make a Hight Availability system with an internal Elastic Load Balancer. Which choice is true regarding internal ELBs in this scenario? (Choose 2 answers)

A.	Internal ELBs should only be launched within private subnets.
B.	Amazon ELB service does not allow subnet selection; instead it will automatically select all the available subnets of the VPC.
C.	Internal ELBs can support only one subnet in each availability zone.
D.	An internal ELB can support all the subnets irrespective of their zones.

Answer: AC
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as elastic load balancers, and EC2 instances. There are two ELBs available with VPC: internet facing and internal (private) ELB. For internal servers, such as App servers the organization can create an internal load balancer in their VPC and then place back-end application instances behind the internal load balancer. The internal load balancer will route requests to the back-end application instances, which are also using private IP addresses and only accept requests from the internal load balancer. The Internal ELB supports only one subnet in each AZ and asks the user to select a subnet while configuring internal ELB.
http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/USVPC_creating_basic_lb.html

QUESTION 353
To ensure failover capabilities on an elastic network interface (ENI), what should you use for incoming traffic?

A.	A Route53 A record
B.	A secondary private IP
C.	A secondary public IP
D.	A secondary ENI

Answer: B
Explanation:
To ensure failover capabilities on an elastic network interface (ENI), consider using a secondary private IP for incoming traffic and if a failure occurs, you can move the interface and/or secondary private IP address to a standby instance.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 354
An organization is setting up a highly scalable application using Elastic Beanstalk. The organization is using ELB and RDS with VPC. The organization has public and private subnets within the cloud.
Which of the below mentioned configurations will not work in this scenario?

A.	To setup RDS in a private subnet and ELB in a public subnet.

B. The configuration must have public and private subnets in the same AZ.
C. The configuration must have two private subnets in separate AZs.
D. The EC2 instance should have a public IP assigned to it.

Answer: D
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. If the organization is planning to implement a scalable secure application using RDS, VPC and ELB the organization should follow below mentioned configurations:
Setup RDS in a private subnet
Setup ELB in a public subnet
Since RDS needs a subnet group, the organization should have two private subnets in the same zone
The ELB needs private and public subnet to be part of same AZs It is not required that instances should have a public IP assigned to them. The instances can be a part of a private subnet and the organization can setup a corresponding routing mechanism.
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/vpc-rds.html

QUESTION 355
In DynamoDB, a projection is _____.

A. systematic transformation of the latitudes and longitudes of the locations inside your table
B. importing data from your file to a table
C. exporting data from a table to your file
D. the set of attributes that is copied from a table into a secondary index

Answer: D
Explanation:
In DynamoDB, a projection is the set of attributes that is copied from a table into a secondary index.
http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GSI.html

QUESTION 356
Which of the following is NOT true of the DynamoDB Console?

A. It allows you to add local secondary indexes to existing tables.
B. It allows you to query a table.
C. It allows you to set up alarms to monitor your table's capacity usage.
D. It allows you to view items stored in a tables, add, update, and delete items.

Answer: A
Explanation:
The DynamoDB Console lets you do the following: Create, update, and delete tables. The throughput calculator provides you with estimates of how many capacity units you will need to request based on the usage information you provide. View items stored in a tables, add, update, and delete items. Query a table. Set up alarms to monitor your table's capacity usage. View your table's top monitoring metrics on real-time graphs from CloudWatch. View alarms configured for each table and create custom alarms.
http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ConsoleDynamoDB.h tml

QUESTION 357
DynamoDB uses _____ only as a transport protocol, not as a storage format.

A.	WDDX
B.	XML
C.	SGML
D.	JSON

Answer: D
Explanation:
DynamoDB uses JSON only as a transport protocol, not as a storage format. The AWS SDKs use JSON to send data to DynamoDB, and DynamoDB responds with JSON, but DynamoDB does not store data persistently in JSON format.
http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Programming.LowLevelAPI.html

QUESTION 358
In DynamoDB, which of the follwoing allows you to set alarms when you reach a specified threshold for a metric?

A.	Alarm Signal
B.	DynamoDB Analyzer
C.	CloudWatch
D.	DynamoDBALARM

Answer: C
Explanation:
CloudWatch allows you to set alarms when you reach a specified threshold for a metric.
http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/MonitoringDynamoDB.html

QUESTION 359
Is it possible to load data from Amazon DynamoDB into Amazon Redshift?

A.	No, you cannot load all the data from DynamoDB table to a Redshift table as it limited by size constraints.
B.	No
C.	No, DynamoDB data types do not correspond directly with those of Amazon Redshift.
D.	Yes

Answer: D
Explanation:
Yes. When you copy data from an Amazon DynamoDB table into Amazon Redshift, you can perform complex data analysis queries on that data. This includes joins with other tables in your Amazon Redshift cluster.
http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/RedshiftforDynamoDB.html

QUESTION 360
In regard to DynamoDB, when you create a table with a hash-and-range key
_____

A.	You must define one or more Local secondary indexes on that table
B.	You must define one or more Global secondary indexes on that table
C.	You can optionally define one or more secondary indexes on that table
D.	You must define one or more secondary indexes on that table

Answer: C
Explanation:
When you create a table with a hash-and-range key, you can optionally define one or more secondary indexes on that table. A secondary index lets you query the data in the table using an alternate key, in addition to queries against the primary key.
http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DataModel.html

QUESTION 361
Someone has recommended a new client to you and you know he is into online gaming and you are almost certain he will want to set up an online gaming site which will require a database service that provides fast and predictable performance with seamless scalability. Which of the following AWS databases would be best suited to an online gaming site?

A.	Amazon SimpleDB
B.	Amazon DynamoDB
C.	Amazon Redshift
D.	Amazon ElastiCache

Answer: B
Explanation:
Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. You can use Amazon DynamoDB to create a database table that can store and retrieve any amount of data, and serve any level of request traffic. Amazon DynamoDB automatically spreads the data and traffic for the table over a sufficient number of servers to handle the request capacity specified by the customer and the amount of data stored, while maintaining consistent and fast performance.
http://aws.amazon.com/documentation/dynamodb/

QUESTION 362
In DynamoDB, "The data is eventually consistent" means that _____.

A.	a read request immediately after a write operation might not show the latest change.
B.	a read request immediately after a write operation shows the latest change.
C.	a write request immediately after a read operation might cause data loss.
D.	a read request immediately after a write operation might cause data loss.

Answer: A
Explanation:
In DynamoDB, it takes time for the update to propagate to all copies. The data is eventually consistent, meaning that a read request immediately after a write operation might not show the latest change.
http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/APISummary.html

QUESTION 363
_____ pricing offers significant savings over the normal price of DynamoDB

provisioned throughput capacity.

A.	Discount Voucher
B.	Reserved Capacity
C.	Discount Service
D.	Reserved Point

Answer: B
Explanation:
Reserved Capacity pricing offers significant savings over the normal price of DynamoDB provisioned throughput capacity. When you buy Reserved Capacity, you pay a one-time upfront fee and commit to paying for a minimum usage level, at the hourly rates indicated above, for the duration of the Reserved Capacity term.
http://aws.amazon.com/dynamodb/pricing/

QUESTION 364
In DynamoDB, which of the following operations is not possible by the console?

A.	Updating an item
B.	Copying an item
C.	Blocking an item
D.	Deleting an item

Answer: C
Explanation:
By using the console to manage DynamoDB, you can perform the following: adding an item, deleting an item, updating an item, and copying an item.
http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AddUpdateDeleteItems.html

QUESTION 365
In order for a table write to succeed, the provisioned throughput settings for the table and global secondary indexes, in DynamoDB, must have _____; otherwise, the write to the table will be throttled.

A.	enough write capacity to accommodate the write
B.	no additional write cost for the index
C.	100 bytes of overhead per index item
D.	the size less than or equal to 1 KB

Answer: A
Explanation:
In order for a table write to succeed in DynamoDB, the provisioned throughput settings for the table and global secondary indexes must have enough write capacity to accommodate the write; otherwise, the write will be throttled.
http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GSI.html

QUESTION 366
In DynamoDB, to get a detailed listing of secondary indexes on a table, you can use the _____ action.

A. BatchGetItem
B. TableName
C. DescribeTable
D. GetItem

Answer: C
Explanation:
In DynamoDB, DescribeTable returns information about the table, including the current status of the table, when it was created, the primary key schema, and any indexes on the table.
http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html

QUESTION 367
In regard to DynamoDB, for which one of the following parameters does Amazon not charge you?

A. Storage cost
B. I/O usage within the same Region
C. Cost per provisioned read units
D. Cost per provisioned write units

Answer: B
Explanation:
In DynamoDB, you will be charged for the storage and the throughput you use rather than for the I/O which has been used.
http://aws.amazon.com/dynamodb/pricing/

QUESTION 368
Complete this statemente: "When you load your table directly from an Amazon _____ table, you have the option to control the amount of provisioned throughput you consume."

A. RDS
B. DataPipeline
C. DynamoDB
D. S3

Answer: C
Explanation:
When you load your table directly from an Amazon DynamoDB table, you have the option to control the amount of Amazon DynamoDB provisioned throughput you consume.
http://docs.aws.amazon.com/redshift/latest/dg/t_Loading_tables_with_the_COPY_command.html

QUESTION 369
Which of the following does Amazon DynamoDB perform?

A. Atomic increment or decrement on scalar values
B. Neither increment nor decrement operations
C. Only increment on vector values
D. Only atomic decrement operations

Answer: A
Explanation:

Amazon DynamoDB allows atomic increment and decrement operations on scalar values.
http://aws.amazon.com/dynamodb/faqs/

QUESTION 370
A Provisioned IOPS volume must be at least _____ GB in size

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

Answer: B
Explanation:
A Provisioned IOPS volume must be at least 10 GB in size
http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Storage.html

QUESTION 371
A user has launched two EBS backed EC2 instances in the US-East-1a region. The user wants to change the zone of one of the instances. How can the user change it?

A.	It is not possible to change the zone of an instance after it is launched
B.	From the AWS EC2 console, select the Actions - > Change zones and specify the new zone
C.	The zone can only be modified using the AWS CLI
D.	Stop one of the instances and change the availability zone

Answer: A
Explanation:
With AWS EC2, when a user is launching an instance he can select the availability zone (AZ) at the time of launch. If the zone is not selected, AWS selects it on behalf of the user. Once the instance is launched, the user cannot change the zone of that instance unless he creates an AMI of that instance and launches a new instance from it.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html

QUESTION 372
In Amazon Elastic Compute Cloud, you can specify storage volumes in addition to the root device volume when you create an AMI or when launching a new instance using _____.

A.	block device mapping
B.	object mapping
C.	batch storage mapping
D.	datacenter mapping

Answer: A
Explanation:
When creating an AMI or launching a new instance, you can assign more than one block storage device to it.
This device will be automatically set ready for you through an automated process known as block device mapping.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/block-device-mapping-concepts.html

QUESTION 373
You create an Amazon Elastic File System (EFS) file system and mount targets for the file system in your Virtual Private Cloud (VPC). Identify the initial permissions you can grant to the group root of your file system.

A.	write-execute-modify
B.	read-execute
C.	read-write-modify
D.	read-write

Answer: B
Explanation:
In Amazon EFS, when a file system and mount targets are created in your VPC, you can mount the remote file system locally on your Amazon Elastic Compute Cloud (EC2) instance. You can grant permissions to the users of your file system. The initial permissions mode allowed for Amazon EFS are:
read-write-execute permissions to the owner root
read-execute permissions to the group root
read-execute permissions to others
http://docs.aws.amazon.com/efs/latest/ug/accessing-fs-nfs-permissions.html

QUESTION 374
You want to mount an Amazon EFS file system on an Amazon EC2 instance using DNS names. Which of the following generic form of a mount target's DNS name must you use to mount the file system?

A.	availability-zone.file-system-id.efs.aws-region.amazonaws.com
B.	efs-system-id.availability-zone.file-aws-region.amazonaws.com
C.	$file-system-id.$availability-zone.$efs.aws-region.$amazonaws.com
D.	#aws-region.#availability-zone.#file-system-id.#efs.#amazonaws.com

Answer: A
Explanation:
An Amazon EFS file system can be mounted on an Amazon EC2 instance using DNS names. This can be done with either a DNS name for the file system or a DNS name for the mount target. To construct the mount target's DNS name, use the following generic form:
availability-zone.file-system-id.efs.aws-region.amazonaws.com
http://docs.aws.amazon.com/efs/latest/ug/mounting-fs.html#mounting-fs-install-nfsclient

QUESTION 375
A user is creating a Provisioned IOPS volume. What is the maximum ratio the user should configure between Provisioned IOPS and the volume size?

A.	30 to 1
B.	50 to 1
C.	10 to 1
D.	20 to 1

Answer: D
Explanation:
Provisioned IOPS SSD (io1) volumes are designed to meet the needs of I/O-intensive workloads, particularly database workloads, that are sensitive to storage performance and consistency. An

io1 volume can range in size from 4 GiB to 16 TiB and you can provision 100 up to 20,000 IOPS per volume. The maximum ratio of provisioned IOPS to requested volume size (in GiB) is 50:1. For example, a 100 GiB volume can be provisioned with up to 5,000 IOPS. Any volume 400 GiB in size or greater allows provisioning up to the 20,000 IOPS maximum.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html

QUESTION 376
Your Application is not highly available, and your on-premises server cannot access the mount target because the Availability Zone (AZ) in which the mount target exists is unavailable. Which of the following actions is recommended?

A.	The application must implement the checkpoint logic and recreate the mount target.
B.	The application must implement the shutdown logic and delete the mount target in the AZ.
C.	The application must implement the delete logic and connect to a different mount target in the same AZ.
D.	The application must implement the restart logic and connect to a mount target in a different AZ.

Answer: D
Explanation:
To make sure that there is continuous availability between your on-premises data center and your Amazon Virtual Private Cloud (VPC), it is suggested that you configure two AWS Direct Connect connections. Your application should implement restart logic and connect to a mount target in a different AZ if your application is not highly available and your on-premises server cannot access the mount target because the AZ in which the mount target exists becomes unavailable.
http://docs.aws.amazon.com/efs/latest/ug/performance.html#performance-onpremises

QUESTION 377
Which of the following Amazon RDS storage types is ideal for applications with light or burst I/O requirements?

A.	Both magnetic and Provisioned IOPS storage
B.	Magnetic storage
C.	Provisioned IOPS storage
D.	None of these

Answer: B
Explanation:
Amazon RDS provides three storage types: magnetic, General Purpose (SSD), and Provisioned IOPS (input/output operations per second). Magnetic (Standard) storage is ideal for applications with light or burst I/O requirements.
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html

QUESTION 378
When I/O performance is more important than fault tolerance, which of the following configurations should be used?

A.	SPAN 10
B.	RAID 1
C.	RAID 0
D.	NFS 1

Answer: C
Explanation:
When I/O performance is more important than fault tolerance, the RAID 0 configuration must be used; for example, as in a heavily used database (where data replication is already set up separately).
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html

QUESTION 379
Amazon Elastic File System (EFS) provides information about the space used for an object by using the space_used attribute of the Network File System Version 4.1 (NFSv4.1). The attribute includes the object's current metered data size and not the metadata size. Which of the following utilities will you use to measure the amount of disk that is used of a file?

A. blkid utility
B. du utility
C. sfdisk utility
D. pydf utility

Answer: B
Explanation:
Amazon EFS reports file system sizes and sizes of objects within a file system. Using the NFSv4.1 space_used attribute for measuring the space used for an object, it reports only the object's current metered data size and not the metadata size. There are two utilities available for measuring disk usage of a file, the du and stat utilities.
https://docs.aws.amazon.com/efs/latest/ug/metered-sizes.html

QUESTION 380
You have custom Network File System (NFS) client settings for your Amazon Elastic File System (EFS). It takes up to three seconds for an Amazon Elastic Compute Cloud (EC2) instance to see a write operation performed on a file system from another Amazon EC2 instance. Which of the following actions should you take to solve the custom NFS settings from causing delays in the write operation?

A. Unmount and remount the file system with the noac option to disable attribute caching.
B. Reduce the number of active users that have files open simultaneously on the instances.
C. Verify that the IP address of the specified mount target is valid.
D. Run the write operation from a different user ID on the same Amazon EC2 instance.

Answer: A
Explanation:
When you set up custom NFS client settings, it takes up to three seconds for an Amazon EC2 instance to see a write operation being performed on a file system from another Amazon EC2 instance. To solve this issue, you must unmount and remount your file system with the noac option to disable attribute caching if the NFS client on the Amazon EC2 instance that is reading the data has attribute caching activated. Attribute cache can also be cleared on demand by using a programming language that is compatible with the NFS procedures. To do this, you must send an ACCESS procedure request immediately before a read request.
http://docs.aws.amazon.com/efs/latest/ug/troubleshooting.html#custom-nfs-settings-write-delays

QUESTION 381
Which of the following rules must be added to a mount target security group to access Amazon Elastic File System (EFS) from an on-premises server?

A.	Configure an NFS proxy between Amazon EFS and the on-premises server to route traffic.
B.	Set up a Point-To-Point Tunneling Protocol Server (PPTP) to allow secure connection.
C.	Permit secure traffic to the Kerberos port 88 from the on-premises server.
D.	Allow inbound traffic to the Network File System (NFS) port (2049) from the on-premises server.

Answer: D
Explanation:
By mounting an Amazon EFS file system on an on-premises server, on-premises data can be migrated into the AWS Cloud. Any one of the mount targets in your VPC can be used as long as the subnet of the mount target is reachable by using the AWS Direct Connect connection. To access Amazon EFS from an on-premises server, a rule must be added to the mount target security group to allow inbound traffic to the NFS port (2049) from the on-premises server.
http://docs.aws.amazon.com/efs/latest/ug/how-it-works.html

QUESTION 382
Which of the following is true of Amazon EBS encryption keys?

A.	Amazon EBS encryption uses the Customer Master Key (CMK) to create an AWS Key Management Service (AWS KMS) master key.
B.	Amazon EBS encryption uses the EBS Magnetic key to create an AWS Key Management Service (AWS KMS) master key.
C.	Amazon EBS encryption uses the EBS Magnetic key to create a Customer Master Key (CMK).
D.	Amazon EBS encryption uses the AWS Key Management Service (AWS KMS) master key to create a Customer Master Key (CMK).

Answer: D
Explanation:
Amazon EBS encryption uses AWS Key Management Service (AWS KMS) master keys when creating encrypted volumes and any snapshots created from your encrypted volumes.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html

QUESTION 383
A user is creating a snapshot of an EBS volume. Which of the below statements is incorrect in relation to the creation of an EBS snapshot?

A.	Its incremental
B.	It is a point in time backup of the EBS volume
C.	It can be used to create an AMI
D.	It is stored in the same AZ as the volume

Answer: D
Explanation:
The EBS snapshots are a point in time backup of the EBS volume. It is an incremental snapshot, but is always specific to the region and never specific to a single AZ. Hence the statement "It is stored in the same AZ as the volume" is incorrect.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSSnapshots.html

QUESTION 384
A user is running a critical batch process which runs for 1 hour and 50 mins every day at a fixed time. Which of the below mentioned options is the right instance type and costing model in this

case if the user performs the same task for the whole year?

A. Instance store backed instance with spot instance pricing.
B. EBS backed instance with standard reserved upfront instance pricing.
C. EBS backed scheduled reserved instance with partial instance pricing.
D. EBS backed instance with on-demand instance pricing.

Answer: D
Explanation:
For Amazon Web Services, the reserved instance (standard or convertible) helps the user save money if the user is going to run the same instance for a longer period. Generally if the user uses the instances around 30-40% of the year annually it is recommended to use RI. Here as the instance runs only for 1 hour 50 minutes daily, or less than 8 percent of the year, it is not recommended to have RI as it will be costlier. At its highest potential savings, you are still paying 25 percent of an annual cost for a reserved instance you are you using less than 2 hours a day, (or less than 8 percent of each year) you are not saving money. Even a scheduled reserved instance has a key limitation, which is that it cannot be stopped or rebooted, only manually terminated with a possibility that it could be restarted. You would have to terminate and restart it within the 1 hour 50 minute window, otherwise you would need to wait until the next day. For a critical daily process, this is likely not an option. Spot Instances are not ideal because the process is critical, and must run for a fixed length of time at a fixed time of day. Spot instances would stop and start based on fluctuations in instance pricing, leaving this process potentially unfinished. The user should use on-demand with EBS in this case. While it has the highest cost, it also has the greatest flexibility to ensure that a critical process like this is always completed.
http://aws.amazon.com/ec2/purchasing-options/reserved-instances/

QUESTION 385
A user has created a launch configuration for Auto Scaling where CloudWatch detailed monitoring is disabled. The user wants to now enable detailed monitoring. How can the user achieve this?

A. Update the Launch config with CLI to set InstanceMonitoringDisabled = false
B. The user should change the Auto Scaling group from the AWS console to enable detailed monitoring
C. Create a new Launch Config with detail monitoring enabled and update the Auto Scaling group
D. Update the Launch config with CLI to set InstanceMonitoring.Enabled = true

Answer: C
Explanation:
CloudWatch is used to monitor AWS as well as the custom services. To enable detailed instance monitoring for a new Auto Scaling group, the user does not need to take any extra steps. When the user creates the Auto Scaling launch config as the first step for creating an Auto Scaling group, each launch configuration contains a flag named InstanceMonitoring.Enabled. The default value of this flag is true. When the user has created a launch configuration with InstanceMonitoring.Enabled = false it will involve multiple steps to enable detail monitoring. The steps are:
. Create a new Launch config with detailed monitoring enabled
. Update the Auto Scaling group with a new launch config
. Enable detail monitoring on each EC2 instance
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/as-metricscollected .html

QUESTION 386

A user is sending a custom metric to CloudWatch. If the call to the CloudWatch APIs has different dimensions, but the same metric name, how will CloudWatch treat all the requests?

A. It will reject the request as there cannot be a separate dimension for a single metric.
B. It will group all the calls into a single call.
C. It will treat each unique combination of dimensions as a separate metric.
D. It will overwrite the previous dimension data with the new dimension data.

Answer: C
Explanation:
A dimension is a key-value pair used to uniquely identify a metric. CloudWatch treats each unique combination of dimensions as a separate metric.
Thus, if the user is making 4 calls with the same metric name but a separate dimension, it will create 4 separate metrics.
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/cloudwatch_concepts.html

QUESTION 387
What is the maximum number of data points for an HTTP data request that a user can include in PutMetricRequest in the CloudWatch?

A. 30
B. 50
C. 10
D. 20

Answer: D
Explanation:
The size of a PutMetricData request of CloudWatch is limited to 8KB for the HTTP GET requests and 40KB for the HTTP POST requests. The user can include a maximum of 20 data points in one PutMetricData request.
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/cloudwatch_concepts.html

QUESTION 388
You have set up a huge amount of network infrastructure in AWS and you now need to think about monitoring all of this. You decide CloudWatch will best fit your needs but you are unsure of the pricing structure and the limitations of CloudWatch. Which of the following statements is TRUE in relation to the limitations of CloudWatch?

A. You get 10 CloudWatch metrics, 10 alarms, 1,000,000 API requests, and 1,000 Amazon SNS email notifications per customer per month for free.
B. You get 100 CloudWatch metrics, 100 alarms, 10,000,000 API requests, and 10,000 Amazon SNS email notifications per customer per month for free.
C. You get 10 CloudWatch metrics, 10 alarms, 1,000 API requests, and 100 Amazon SNS email notifications per customer per month for free.
D. You get 100 CloudWatch metrics, 100 alarms, 1,000,000 API requests, and 1,000 Amazon SNS email notifications per customer per month for free.

Answer: A
Explanation:
Amazon CloudWatch monitors your Amazon Web Services (AWS) resources and the applications

you run on AWS in real-time. You can use CloudWatch to collect and track metrics, which are the variables you want to measure for your resources and applications.
CloudWatch has the following limits:
You get 10 CloudWatch metrics, 10 alarms, 1,000,000 API requests, and 1,000 Amazon SNS email notifications per customer per month for free.
You can assign up to 10 dimensions per metric.
You can create up to 5000 alarms per AWS account.
Metric data is kept for 2 weeks.
The size of a PutMetricData request is limited to 8KB for HTTP GET requests and 40KB for HTTP POST requests.
You can include a maximum of 20 MetricDatum items in one PutMetricData request. A MetricDatum can contain a single value or a StatisticSet representing many values.
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/cloudwatch_limits.html

QUESTION 389
A user is trying to send custom metrics to CloudWatch using the PutMetricData APIs. Which of the below mentioned points should the user needs to take care while sending the data to CloudWatch?

A.	The size of a request is limited to 8KB for HTTP GET requests and 40KB for HTTP POST requests
B.	The size of a request is limited to 16KB for HTTP GET requests and 80KB for HTTP POST requests
C.	The size of a request is limited to 128KB for HTTP GET requests and 64KB for HTTP POST requests
D.	The size of a request is limited to 40KB for HTTP GET requests and 8KB for HTTP POST requests

Answer: A
Explanation:
With AWS CloudWatch, the user can publish data points for a metric that share not only the same time stamp, but also the same namespace and dimensions. CloudWatch can accept multiple data points in the same PutMetricData call with the same time stamp. The only thing that the user needs to take care of is that the size of a PutMetricData request is limited to 8KB for HTTP GET requests and 40KB for HTTP POST requests.
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/cloudwatch_concepts.html

QUESTION 390
You set up your first Lambda function and want to set up some Cloudwatch metrics to monitor your function. Which of the following Lambda metrics can Cloudwatch monitor?

A.	Total requests only
B.	Status Check Failed, total requests, and error rates
C.	Total requests and CPU utilization
D.	Total invocations, errors, duration, and throttles

Answer: D
Explanation:
AWS Lambda automatically monitors functions on your behalf, reporting metrics through Amazon CloudWatch (CloudWatch). These metrics include total invocations, errors, duration, and throttles.
http://docs.aws.amazon.com/lambda/latest/dg/monitoring-functions-metrics.html

QUESTION 391
How many metrics are supported by CloudWatch for Auto Scaling ?

A.	7 metrics and 5 dimension
B.	5 metrics and 1 dimension
C.	1 metric and 5 dimensions
D.	8 metrics and 1 dimension

Answer: D
Explanation:
AWS Auto Scaling supports both detailed as well as basic monitoring of the CloudWatch metrics. Basic monitoring happens every 5 minutes, while detailed monitoring happens every minute. It supports 8 metrics and 1 dimension.
The metrics are:
GroupMinSize
GroupMaxSize
GroupDesiredCapacity
GroupInServiceInstances
GroupPendingInstances
GroupStandbyInstances
GroupTerminatingInstances
GroupTotalInstances
The dimension is AutoScalingGroupName
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/supported_services.html

QUESTION 392
A user has enabled detailed CloudWatch monitoring with the AWS Simple Notification Service. Which of the below mentioned statements helps the user understand detailed monitoring better?

A.	SNS cannot provide data every minute
B.	SNS will send data every minute after configuration
C.	There is no need to enable since SNS provides data every minute
D.	AWS CloudWatch does not support monitoring for SNS

Answer: A
Explanation:
CloudWatch is used to monitor AWS as well as the custom services. It provides either basic or detailed monitoring for the supported AWS products. In basic monitoring, a service sends data points to CloudWatch every five minutes, while in detailed monitoring a service sends data points to CloudWatch every minute. The AWS SNS service sends data every 5 minutes. Thus, it supports only the basic monitoring. The user cannot enable detailed monitoring with SNS.
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/supported_services.html

QUESTION 393
An AWS account owner has setup multiple IAM users. One of these IAM users, named John, has CloudWatch access, but no access to EC2 services. John has setup an alarm action which stops EC2 instances when their CPU utilization is below the threshold limit. When an EC2 instance's CPU Utilization rate drops below the threshold John has set, what will happen and why?

A.	CloudWatch will stop the instance when the action is executed

B. Nothing will happen. John cannot set an alarm on EC2 since he does not have the permission.
C. Nothing will happen. John can setup the action, but it will not be executed because he does not have EC2 access through IAM policies.
D. Nothing will happen because it is not possible to stop the instance using the CloudWatch alarm

Answer: C
Explanation:
Amazon CloudWatch alarms watch a single metric over a time period that the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The user can setup an action which stops the instances when their CPU utilization is below a certain threshold for a certain period of time. The EC2 action can either terminate or stop the instance as part of the EC2 action. If the IAM user has read/write permissions for Amazon CloudWatch but not for Amazon EC2, he can still create an alarm. However, the stop or terminate actions will not be performed on the Amazon EC2 instance.
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/UsingAlarmActions.html

QUESTION 394
A user has configured an EC2 instance in the US-East-1a zone. The user has enabled detailed monitoring of the instance. The user is trying to get the data from CloudWatch using a CLI. Which of the below mentioned CloudWatch endpoint URLs should the user use?

A. monitoring.us-east-1a.amazonaws.com
B. cloudwatch.us-east-1a.amazonaws.com
C. monitoring.us-east-1.amazonaws.com
D. monitoring.us-east-1-a.amazonaws.com

Answer: C
Explanation:
The CloudWatch resources are always region specific and they will have the end point as region specific. If the user is trying to access the metric in the US-East-1 region, the endpoint URL will be: monitoring.us-east-1.amazonaws.com
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/regions_endpoints.html

QUESTION 395
Which of the following is not included in the metrics sent from Billing to Amazon CloudWatch?

A. Recurring fees for AWS products and services
B. Total AWS charges
C. One-time charges and refunds
D. Usage charges for AWS products and services

Answer: C
Explanation:
Usage charges and recurring fees for AWS products and services are included in the metrics sent from Billing to Amazon CloudWatch.
You will have a metric for total AWS charges, as well as one additional metric for each AWS product or service that you use.
However, one-time charges and refunds are not included.
https://aws.amazon.com/blogs/aws/monitor-estimated-costs-using-amazon-cloudwatch-billing-metrics-and-alarms/

QUESTION 396
After your Lambda fuction has been running for some time, you need to look at some metrics to ascertain how your function is performing and decide to use the AWS CLI to do this. Which of the following commands must be used to access these metrics using the AWS CLI?

A.	mon-list-metrics and mon-get-stats
B.	list-metrics and get-metric-statistics
C.	ListMetrics and GetMetricStatistics
D.	list-metrics and mon-get-stats

Answer: B
Explanation:
AWS Lambda automatically monitors functions on your behalf, reporting metrics through Amazon CloudWatch.
To access metrics using the AWS CLI
Use the list-metrics and get-metric-statistics commands.
http://docs.aws.amazon.com/lambda/latest/dg/monitoring-functions-access-metrics.html

QUESTION 397
In Amazon CloudWatch, you can publish your own metrics with the put-metric-data command. When you create a new metric using the put-metric-data command, it can take up to two minutes before you can retrieve statistics on the new metric using the get-metric-statistics command. How long does it take before the new metric appears in the list of metrics retrieved using the list-metrics command?

A.	After 2 minutes
B.	Up to 15 minutes
C.	More than an hour
D.	Within a minute

Answer: B
Explanation:
You can publish your own metrics to CloudWatch with the put-metric-data command (or its Query API equivalent PutMetricData). When you create a new metric using the put-metric-data command, it can take up to two minutes before you can retrieve statistics on the new metric using the get-metric-statistics command. However, it can take up to fifteen minutes before the new metric appears in the list of metrics retrieved using the list-metrics command.
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/publishingMetrics.html

QUESTION 398
A user is running a batch process on EBS backed EC2 instances. The batch process launches few EC2 instances to process hadoop Map reduce jobs which can run between 50 ?600 minutes or sometimes for even more time. The user wants a configuration that can terminate the instance only when the process is completed. How can the user configure this with CloudWatch?

A.	Configure a job which terminates all instances after 600 minutes
B.	It is not possible to terminate instances automatically
C.	Configure the CloudWatch action to terminate the instance when the CPU utilization falls below 5%
D.	Set up the CloudWatch with Auto Scaling to terminate all the instances

Answer: C

Explanation:
Amazon CloudWatch alarm watches a single metric over a time period that the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The user can setup an action which terminates the instances when their CPU utilization is below a certain threshold for a certain period of time. The EC2 action can either terminate or stop the instance as part of the EC2 action.
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/UsingAlarmActions.html

QUESTION 399
A user is trying to understand the detailed CloudWatch monitoring concept. Which of the below mentioned services does not provide detailed monitoring with CloudWatch?

A.	AWS RDS
B.	AWS ELB
C.	AWS Route53
D.	AWS EMR

Answer: D
Explanation:
CloudWatch is used to monitor AWS as well as the custom services. It provides either basic or detailed monitoring for the supported AWS products. In basic monitoring, a service sends data points to CloudWatch every five minutes, while in detailed monitoring a service sends data points to CloudWatch every minute. Services, such as RDS, EC2, Auto Scaling, ELB, and Route 53 can provide the monitoring data every minute.
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/supported_services.html