Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

https://www.2passeasy.com/dumps/AWS-Certified-Solutions-Architect-Professional/

NEW QUESTION 1
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 ofthe 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.
Reference:
http://docs.aws.amazon.com/AmazonS3/latest/dev/access-control-auth-workflow-object-operation.htmI

NEW QUESTION 2
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 fonrvard 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. Reference:
http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html

NEW QUESTION 3
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.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Appendix_Limits.html

NEW QUESTION 4
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 ThanksgMng 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.
Reference: http://media.amazonwebservices.com/AWS_Web_Hosting_Best_Practices.pdf

NEW QUESTION 5
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 indMdual IAM use

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-I7",
"Statement":
[{ "Effect": "AIlow", "Action": "ec2:RunInstances", "Resource":
["arn:aws:ec2:region::image/ami-*", "arn:aws:ec2:region:accountzsubnet/subnet-1a2b3c4d", "arn:aws:ec2:region:account:network-interface/*", "arn:aws:ec2:region:account:vo|ume/*", "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
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_IANI.htmI

NEW QUESTION 6
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)). Reference: https://aws.amazon.com/premiumsupport/faqs/

NEW QUESTION 7
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.
Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_|imits.htmI#limits_ec2

NEW QUESTION 8
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.
Reference: http://docs.aws.amazon.com/cognito/devguide/identity/concepts/authentication-f|ow/

NEW QUESTION 9
In Amazon EIastiCache, 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 EIastiCache 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 capacit

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.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/CacheNode.Memcached.html

NEW QUESTION 10
IV|apMySite 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 subne
B. Launch an application instance in the public subnet while RDS is launched in the private subnet.
C. Setup a public and two private subnets in different AZs within a VPC and create a subnet grou
D. Launch RDS with that subnet group.
E. Create a network interface and attach two subnets to i
F. Attach that network interface with RDS while launching a DB instance.
G. 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.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.html

NEW QUESTION 10
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 actMty that uses the resources is running
C. AWS Data Pipeline terminates AWS Data Pipeline-managed compute resources every 12 hours.
D. When the final actMty that uses the resources has completed successfully orfailed

Answer: D

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

NEW QUESTION 14
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. "a||ow"
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.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPoIicyLanguage_EvaluationLogic.html

NEW QUESTION 18
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-I7",
"Statement": [{
"Effect": "A||ow", "Action": ["iam:AddUserToGroup",
"iam:RemoveUserFromGroup", "iam:GetGroup"
]!
"Resource": "arn:aws:iam:: I23412341234:group/TestingGroup"
}I

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-I7",
"Statement": [{
"Effect": "A||ow", "Action": ["iam:AddUserToGroup",
"iam:RemoveUserFromGroup", "iam:GetGroup"
]!
"Resource": "arn:aws:iam:: 123412341234:group/ TestingGroup "
}I
Reference:
http://docs.aws.amazon.com/IAM/latest/UserGuide/CredentiaIs-Permissions-examples.htm|#creds-po|ici es-credentials

NEW QUESTION 21
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 connectMty) and the instance type EBS dedicated bandwidth exceeds the IOPS more than he has provisioned.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-io-characteristics.html

NEW QUESTION 22
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:|nvokeAsync
D. sts:AssumeRoIe

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:AssumeRoIe action.
Reference: http://docs.aws.amazon.com/|AM/Iatest/UserGuide/id_ro|es_manage_modify.html

NEW QUESTION 24
One of your AWS Data Pipeline actMties 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 actMty fails if all of its actMty attempts return with a failed state. By default, an actMty 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.
Reference: https://aws.amazon.com/datapipe|ine/faqs/

NEW QUESTION 28
True or False: In Amazon EIastiCache 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 tunin

Answer: C

Explanation:
In Amazon EIastiCache, 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.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/Replication.Redis.Groups.html

NEW QUESTION 33
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.
Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#|imits_ec2

NEW QUESTION 34
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 environmen

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 Nlemcached 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.
Reference: https://aws.amazon.com/elasticache/

NEW QUESTION 35
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 cost

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 connectMty 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.
Reference: http://aws.amazon.com/directconnect/#details

NEW QUESTION 36
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 dat
C. Thus, whenever the instance is booted the script assigns the fixed MAC address to that instance.
D. The instance MAC address never change
E. Thus, it is not required to register the MAC address every time.
F. 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.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AESDG-chapter-instancedata.html

NEW QUESTION 40
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 sewers 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 ligh

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.
Reference: http://d36cz9buwru1tt.cloudfront.net/AWS_Disaster_Recovery.pdf

NEW QUESTION 44
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. Reference: http://aws.amazon.com/e|asticache/faqs/

NEW QUESTION 46
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.
Reference: http://docs.aws.amazon.com/datapipeline/latest/DeveIoperGuide/what-is-datapipeline.htmI

NEW QUESTION 49
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.
Reference:
http://docs.rackspace.com/Ioadbalancers/api/v1.0/clb-devguide/content/Concepts-d1e233.htmI

NEW QUESTION 51
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 identitie

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. Reference: http://docs.aws.amazon.com/cognito/devguide/identity/identity-pool|s/

NEW QUESTION 54
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": "AIlow", "Action": ["Describe"], "Resource": "BiIling"
B. "Effect": "AIIow", "Action": ["aws-portal: ViewBi||ing"], "Resource": "*"
C. "Effect": "AIIow", "Action": ["aws-portal:ViewUsage"], "Resource": "*"
D. "Effect": "AIIow", "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": "A||ow", "Action": [
"aws-portal:ViewUsage"
]!
"Resource": "*"
} I
}
Reference: http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/billing-permissions-ref.html

NEW QUESTION 55
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.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Appendix_Limits.html

NEW QUESTION 59
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.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-io-characteristics.html

NEW QUESTION 62
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.
Reference: http://docs.aws.amazon.com/amazonglacier/latest/dev/configuring-notifications-console.html

NEW QUESTION 63
ExamKiIIer 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 ExamKil|er 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.
Reference:
http://docs.aws.amazon.com/AmazonVPC/Iatest/PeeringGuide/peering-configurations-full-access.htmI#t hree-vpcs-full-access

NEW QUESTION 64
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.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html

NEW QUESTION 66
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.
Reference: http://docs.aws.amazon.com/IANI/latest/UserGuide/LimitationsOnEntities.html

NEW QUESTION 70
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.
Reference: http://docs.aws.amazon.com/IAM/Iatest/UserGuide/reference_policies_eIements.html

NEW QUESTION 74
Which of the following cache engines does Amazon EIastiCache support?

A. Amazon EIastiCache supports Memcached and Redis.
B. Amazon EIastiCache supports Redis and WinCache.
C. Amazon EIastiCache supports Memcached and Hazelcast.
D. Amazon EIastiCache supports Memcached onl

Answer: A

Explanation:
The cache engines supported by Amazon EIastiCache are Memcached and Redis.
Reference: http://docs.aws.amazon.com/AmazonEIastiCache/latest/UserGuide/SeIectEngine.html

NEW QUESTION 76
You have been given the task to define multiple AWS Data Pipeline schedules for different actMties 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 actMty via its schedule field
C. Defining multiple schedule objects in your pipeline definition file and associating the desired schedule to the correct actMty via its schedule field
D. Defining multiple schedule objects in the schedule field

Answer: C

Explanation:
To define multiple schedules for different actMties 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 actMty 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. Reference: https://aws.amazon.com/datapipeline/faqs/

NEW QUESTION 81
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 creatio

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

NEW QUESTION 84
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.
Reference: http://media.amazonwebservices.com/pdf/AWS_Security_Whitepaper.pdf

NEW QUESTION 88
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.
Reference: http://aws.amazon.com/ec2/purchasing-options/spot-instances/

NEW QUESTION 92
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.
Reference: http://docs.aws.amazon.com/sns/latest/dg/SNSMobi|ePushPrereq.htmI

NEW QUESTION 97
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 addres

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.
Reference:
http://docs.aws.amazon.com/AmazonVPC/Iatest/UserGuide/VPC_NAT_|nstance.htm|#EIP_Disab|e_Src DestCheck

NEW QUESTION 100
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.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/NlultipleIP.htmI

NEW QUESTION 104
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. Nlongo 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. Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVo|umeTypes.htm|

NEW QUESTION 109
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.

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

NEW QUESTION 112
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. StringEquaIsIgnoreCase
B. StringNotEquaIsIgnoreCase
C. StringLikeStringEqua|s
D. StringNotEqua|s

Answer: A

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

NEW QUESTION 115
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.
Reference:
http://docs.aws.amazon.com/datapipeline/latest/DeveIoperGuide/dp-how-tasks-scheduled.htmI

NEW QUESTION 119
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 Nlanagement (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 resourc

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

NEW QUESTION 120
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'II need to implement that functionality in a custom task runner.
Reference:
http://docs.aws.amazon.com/datapipeline/latest/DeveIoperGuide/dp-how-remote-taskrunner-client.html

NEW QUESTION 123
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.
Reference: http://docs.aws.amazon.com/Iambda/latest/dg/lambda-dg.pdf

NEW QUESTION 125
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 subne

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. Reference:
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Introduction.htm|

NEW QUESTION 129
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 subnet
B. Both the subnets can be in the same or separate AZ.
C. The organization should keep minimum of one IP address in each subnet reserved for RDS failover.
D. If the organization is connecting RDS from the internet it must enable the VPC attributes DNS hostnames and DNS resolution.
E. 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.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.html

NEW QUESTION 131
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 interfac

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.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AESDG-chapter-instancedata.html

NEW QUESTION 134
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.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-peering.htmI

NEW QUESTION 139
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. Reference: http://aws.amazon.com/about-aws/whats-new/2012/12/20/amazon-elasticache-announces-support-for-a mazon-vpc/

NEW QUESTION 141
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.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html

NEW QUESTION 146
In the context of Amazon E|astiCache 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). Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/ECEvents.Viewing.html

NEW QUESTION 147
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 theusen
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 usen
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.
Reference: http://aws.amazon.com/cognito/faqs/

NEW QUESTION 149

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.
Reference: http://awsdocs.s3.amazonaws.com/SQS/2011-10-01/sqs-dg-2011-10-01.pdf

NEW QUESTION 153
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.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/vpc-migrate.html

NEW QUESTION 154
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. Reference:
http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html

NEW QUESTION 158
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.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/WhatIs.html

NEW QUESTION 159
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 sewer or instance on which your application and the HSNI client is running must have network (IP) reachability to the HSNI.
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 instanc

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 connectMty in a variety of ways, including operating your application in the same VPC, with VPC peering, with a VPN connection, or with Direct Connect.
Reference: https://aws.amazon.com/cloudhsm/faqs/

NEW QUESTION 162
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.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/CacheSecurityGroup.html

NEW QUESTION 166
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 sensitMty 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.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-workload-demand.htmI

NEW QUESTION 170
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 subne

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.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html#VPCSubnet

NEW QUESTION 171
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 networ

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.

NEW QUESTION 173
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 AW

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.
Reference: http://docs.aws.amazon.com/AmazonC|oudFront/latest/DeveloperGuide/Troubleshooting.html

NEW QUESTION 177
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.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_EvaluationLogic.html

NEW QUESTION 180
Which of the following AWS services can be used to define alarms to trigger on a certain actMty, such as actMty 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 actMties such as success, failure, or delay by creating an alarm object and referencing it in the onFail, onSuccess, or onLate slots of the actMty object.
Reference: https://aws.amazon.com/datapipe|ine/faqs/

NEW QUESTION 182
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 instance

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

NEW QUESTION 185
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. Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/WhatIs.NIanaging.html

NEW QUESTION 188
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 locatio

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

NEW QUESTION 191
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 grou
B. 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.
C. In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 50 rules to a security grou
D. 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.
E. In EC2-Classic, you can associate an instance with up to 5 security groups and add up to 100 rules to a security grou
F. 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.
G. In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 100 rules to a security grou
H. 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.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html

NEW QUESTION 196
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. Reference: http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/CreateVPCForELB.html

NEW QUESTION 199
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 Connec

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

NEW QUESTION 204
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 IV|anagement 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: D

NEW QUESTION 206
You are looking to migrate your Development (Dev) and Test environments to AWS. You have decided to use separate AWS accounts to host each environment. You plan to link each accounts bill to a Master AWS account using Consolidated Billing. To make sure you Keep within budget you would like to implement a way for administrators in the Master account to have access to stop, delete and/or terminate resources in both the Dev and Test accounts. Identify which option will allow you to achieve this goal.

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

Answer: C

NEW QUESTION 210
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. Configure your backup software to use S3 as the target for your data backups.
C. Configure your backup software to use Glacier as the target for your data backups.
D. Use Storage Gateway and configure it to use Gateway Stored volume

Answer: A

NEW QUESTION 215
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. Dunning a DR test you notice that when you disable all web sewers in one of the regions Route53 does not automatically direct all users to the other region. What could be happening? (Choose 2 answers)

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

Answer: BE

NEW QUESTION 219
Your startup wants to implement an order fulfillment process for selling a personalized gadget that needs an average of 3-4 days to produce with some orders taking up to 6 months you expect 10 orders per day on your first day. 1000 orders per day after 6 months and 10,000 orders after 12 months.
Orders coming in are checked for consistency men dispatched to your manufacturing plant for production quality control packaging shipment and payment processing If the product does not meet the quality standards at any stage of the process employees may force the process to repeat a step Customers are notified via email about order status and any critical issues with their orders such as payment failure.
Your case architecture includes AWS Elastic Beanstalk for your website with an RDS MySQL instance for customer data and orders.
How can you implement the order fulfillment process while making sure that the emails are delivered reliably?

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

Answer: C

NEW QUESTION 220
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 autoscaimg 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 RDS with read replicas.
C. Stateful instances for the web and application tier in an autoscaling group monitored with CloudWatc
D. And multi-AZ RDS.
E. 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

NEW QUESTION 225
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 indMdual 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 c3.8x|arge instance in us-east-1. Provision 4x1TB EBS volumes, attach them to the instance, and configure them as a single RAID 5 volum
B. Ensure that EBS snapshots are performed every 15 minutes.
C. Instantiate a c3.8xlarge instance in us-east-1. Provision 3xITB EBS volumes, attach them to the Instance, and configure them as a single RAID 0 volum
D. Ensure that EBS snapshots are performed every 15 minutes.
E. Instantiate an i2.8xlarge instance in us-east-1
F. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instanc
G. Provision 3x1TB EBS volumes, attach them to the instance, and configure them as a second RAID 0 volum
H. Configure synchronous, block-level replication from the ephemeral-backed volume to the EBS-backed volume.
I. Instantiate a c3.8xlarge instance in us-east-1. Provision an AWS Storage Gateway and configure it for 3 TB of storage and 100,000 IOP
J. Attach the volume to the instance.
K. Instantiate an i2.8xlarge instance in us-east-1
L. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instanc
M. Configure synchronous, blocklevel replication to an identically configured instance in us-east-1b.

Answer: C

NEW QUESTION 228
A large real-estate brokerage is exploring the option o(adding a cost-effective location based alert to their existing mobile application The application backend infrastructure currently runs on AWS Users who opt in to this service will receive alerts on their mobile device regarding real-estate otters in proximity to their location. For the alerts to be relevant delivery time needs to be in the low minute count the existing mobile app has 5 million users across the US. Which one of the following architectural suggestions would you make to the customer?

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

NEW QUESTION 232
Your department creates regular analytics reports from your company's log files All log data is collected in Amazon S3 and processed by daily Amazon Elastic MapReduce (EMR) jobs that generate daily PDF reports and aggregated tables in CSV format for an Amazon Redshift data warehouse.
Your CFO requests that you optimize the cost structure for this system.
Which of the following alternatives will lower costs without compromising average performance of the system or data integrity for the raw data?

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

Answer: C

NEW QUESTION 237
You require the ability to analyze a large amount of data, which is stored on Amazon S3 using Amazon Elastic Map Reduce. You are using the cc2 8x large 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 more smaller flies on Amazon S3.
B. Add additional cc2 8x large instances by introducing a task group.
C. Use smaller instances that have higher aggregate I/O performance.
D. Create fewer, larger files on Amazon S3.

Answer: C

NEW QUESTION 240
Your customer wishes to deploy an enterprise application to AWS which will consist of several web servers, several application servers and a small (50GB) Oracle database information is stored, both in the database and the file systems of the various servers. The backup system must support database recovery whole server and whole disk restores, and indMdual file restores with a recovery time of no more than two hours. They have chosen to use RDS Oracle as the database
Which backup architecture will meet these requirements?

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

Answer: A

NEW QUESTION 244
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

A. The logistic software has a 3-tierarchitecture and currently uses MySQL 5.6 for data persistenc
B. 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'?
C. For each regional deployment, use RDS MySQL with a master in the region and a read replica in theHQ region
D. For each regional deployment, use NIySQL on EC2 with a master in the region and send hourly EBS snapshots to the HQ region
E. For each regional deployment, use RDS MySQL with a master in the region and send hourly RDS snapshots to the HQ region
F. 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
G. Use Direct Connect to connect all regional MySQL deployments to the HQ region and reduce network latency for the batch process

Answer: A

NEW QUESTION 245
An enterprise wants to use a third-party SaaS application. The SaaS application needs to have access to issue several API commands to discover Amazon EC2 resources running within the enterprise's account The enterprise has internal security policies that require any outside access to their environment must conform to the principles of least prMlege and there must be controls in place to ensure that the credentials used by the SaaS vendor cannot be used by any other third party. Which of the following would meet all of these conditions?

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

Answer: C

NEW QUESTION 246
Company B is launching a new game app for mobile devices. Users will log into the game using their existing social media account to streamline data capture. Company B would like to directly save player data and scoring information from the mobile app to a DynamoDS table named Score Data When a user saves their game the progress data will be stored to the Game state S3 bucket. What is the best approach for storing data to DynamoDB and S3?

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

Answer: B

NEW QUESTION 249
Your company is getting ready to do a major public announcement of a social media site on AWS. The website is running on EC2 instances deployed across multiple Availability Zones with a MuIti-AZ RDS MySQL Extra Large DB Instance. The site performs a high number of small reads and writes per second and relies on an eventual consistency model. After comprehensive tests you discover that there is read contention on RDS MySQL. Which are the best approaches to meet these requirements? (Choose 2 answers)

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

Answer: AC

NEW QUESTION 253
You are designing an intrusion detection prevention (IDS/IPS) solution for a customer web application in a single VPC. You are considering the options for implementing IOS IPS protection for traffic coming from the Internet.
Which of the following options would you consider? (Choose 2 answers)

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

Answer: BD

NEW QUESTION 257
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 Cloud Watch monitors the number of Job requests (queued messages) and an Auto Scaling group adds or deletes batch sewers automatically based on parameters set in Cloud Watch alarms. You can use this architecture to implement which of the following features in a cost effective and efficient manner?

A. Reduce the overall lime 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.
B. Implement fault tolerance against EC2 instance failure since messages would remain in SQS and worn can continue with recovery of EC2 instances implement fault tolerance against SQS failure by backing up messages to S3.
C. Implement message passing between EC2 instances within a batch by exchanging messages throughSQS.
D. Coordinate number of EC2 instances with number of job requests automatically thus Improving cost effectiveness.
E. Handle high priority jobs before lower priority jobs by assigning a priority metadata field to SQS messages.

Answer: D

NEW QUESTION 261
An International company has deployed a multi-tier web application that relies on DynamoDB in a single region For regulatory reasons they need disaster recovery capability In a separate region with a Recovery Time Objective of 2 hours and a Recovery Point Objective of 24 hours. They should synchronize their data on a regular basis and be able to provision me web application rapidly using CIoudFormation.
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 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.
C. 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.
D. Send also each Ante into an SQS queue in me second region; use an auto-scaling group behind the SQS queue to replay the write in the second region.

Answer: A

NEW QUESTION 262
A web company is looking to implement an external payment service into their highly available application deployed in a VPC Their application EC2 instances are behind a public lacing ELB Auto scaling is used to add additional instances as traffic increases under normal load the application runs 2 instances in the Auto Scaling group but at peak it can scale 3x in size. The application instances need to communicate with the payment service over the Internet which requires whitelisting of all public IP addresses used to communicate with it. A maximum of 4 whitelisting IP addresses are allowed at a time and can be added through an API.
How should they architect their solution?

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

Answer: D

NEW QUESTION 267
Your website is serving on-demand training videos to your workforce. Videos are uploaded monthly in high resolution MP4 format. Your workforce is distributed globally often on the move and using company-provided tablets that require the HTTP Live Streaming (HLS) protocol to watch a video. Your company has no video transcoding expertise and it required you may need to pay for a consultant.
How do you implement the most cost-efficient architecture without compromising high availability and quality of video delivery'?

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

L. CloudFront to serve HLS transcoded videos from Glacier.

Answer: C

NEW QUESTION 269
A customer has established an AWS Direct Connect connection to AWS. The link is up and routes are being advertised from the customer's end, however the customer is unable to connect from EC2 instances inside its VPC to servers residing in its datacenter.
Which of the following options provide a viable solution to remedy this situation? (Choose 2 answers)

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

Answer: AC

NEW QUESTION 271
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 fileet 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 c3.8xIarge instances; those instances can potentially yield a network throughput 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 ap
C. Implement client-side logic to directly upload the file to Amazon S3 using the given credentials and S3 prefix.
D. Re-architect your ingest pattern, and move your web application instances into a VPC public subne
E. 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 apprequests; this approach will significantly reduce the cost by bypassing Elastic Load Balancing.
F. 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 ap
G. 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

NEW QUESTION 274
You require the ability to analyze a customer's clickstream data on a website so they can do behavioral analysis. Your customer needs to know what sequence of pages and ads their customer clicked on. This data will be used in real time to modify the page layouts as customers click through the site to increase stickiness and advertising click-through. Which option meets the requirements for captioning and analyzing this data?

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

Answer: B

NEW QUESTION 275
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 sewers, two database sewers and one NAT instance tor a total of seven EC2 instances The web. Application and database sewers 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 (raffle gradually increases in the first few days following the deployment, so you attempt to double the number of instances in each tier of the application to handle the new load unfortunately some of these new instances fail to launch.
Which of the following could be the root caused? (Choose 2 answers)

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

Answer: CE

NEW QUESTION 279
Your company produces customer commissioned one-of-a-kind skiing helmets combining nigh fashion with custom technical enhancements Customers can show off their IndMduality on the ski slopes and have access to head-up-displays. GPS rear-view cams and any other technical innovation they wish to embed in the helmet.
The current manufacturing process is data rich and complex including assessments to ensure that the custom electronics and materials used to assemble the helmets are to the highest standards Assessments are a mixture of human and automated assessments you need to add a new set of assessment to model the failure modes of the custom electronics using GPUs with CUDA, across a cluster of servers with low latency networking.
What architecture would allow you to automate the existing process using a hybrid approach and ensure that the architecture can support the evolution of

processes over time?

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

Answer: B

NEW QUESTION 283
You are designing an SSUTLS solution that requires HTTPS clients to be authenticated by the Webserver 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 ELB with TCP listeners on TCP/443. And place the Web servers behind it.
B. Configure your Web servers with EIP
C. 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 sewers behind it.
E. Configure your web sewers as the origins for a CIoudFront distributio
F. Use custom SSL certificates on your C|oudFront distribution.

Answer: AB

NEW QUESTION 285
A company is building a voting system for a popular TV show, viewers win 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 their Amazon.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 can handle the rapid influx of traffic while maintaining good performance but also wants to keep costs to a minimum. Which of the design patterns below should they use?

A. Use CIoudFront 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.
B. Use CIoudFront 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 tableto store the users vote.
C. Use CIoudFront 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.
D. Use CIoudFront 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 sewers win process the users vote and store the result into an SQS queue using IAM Roles for EC2 Instances to gain permissions to the SQS queu
E. A set of application sewers will then retrieve the items from the queue and store the result into a DynamoDB table.

Answer: D

NEW QUESTION 289
Which is a valid Amazon Resource name (ARN) for IAM?

A. aws:iam::123456789012:instance-profile/\Nebserver
B. arn:aws:iam::123456789012:instance-profile/Webserver
C. 123456789012:aws:iam::instance-profi|e/\Nebserver
D. arn:aws:iam::123456789012::instance-profile/\Nebserver

Answer: B

NEW QUESTION 290
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 RDS with Redshift for the oaten analysis and SQS to send a message to the on-premises system to update the dashboard
C. Create an RDS Read Replica for the batch analysis and SNS to notify me on-premises system to update the dashboard
D. Create an RDS Read Replica for the batch analysis and SQS to send a message to the on-premises system to update the dashboard.

Answer: A

NEW QUESTION 294
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. 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.

B. Generate the reports by querying the synchronously replicated standby RDS NIySQL instance maintained through Nlulti-AZ.
C. Launch a RDS Read Replica connected to your Multi AZ master database and generate reports by querying the Read Replica.
D. Generate the reports by querying the ElastiCache database caching tie

Answer: C

NEW QUESTION 296
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 IANI user, and retrieve the IAM user's credentials from the EC2 instance user data.

Answer: B

NEW QUESTION 301
A 3-tier e-commerce web application is current deployed on-premises and will be migrated to AWS for greater scalability and elasticity The web server currently shares read-only data using a network distributed file system The app server tier uses a clustering mechanism for discovery and shared session state that depends on IP multicast The database tier uses shared-storage clustering to provide database fall over capability, and uses several read slaves for scaling Data on all servers and the distributed file system directory is backed up weekly to off-site tapes
Which AWS storage and database architecture meets the requirements of the application?

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

Answer: C

NEW QUESTION 304
Your company plans to host a large donation website on Amazon Web Sewices (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 sewice should you use?

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

Answer: B

NEW QUESTION 309
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 sewers: subnet-258bc44d
Application servers: subnet-248bc44c Database sewers: subnet-9189c6f9 Route Tables:
rrb-218bc449 rtb-238bc44b Associations:
subnet-258bc44d : rtb-218bc449 subnet-248bc44c : rtb-238bc44b subnet-9189c6f9 : rtb-238bc44b
You are now ready to begin deploying EC2 instances into the VPC Web servers must have direct access to the internet Application and database servers cannot have direct access to the internet.
Which configuration below will allow you the ability to remotely administer your application and database servers, as well as allow these sewers 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 subneb258bc44d.
D. Create a bastion and NAT instance in subnet-258bc44d, add a route from rtb-238bc44b toIgw-2d8bc445, and a new NACL that allows access between subnet-258bc44d and subnet-248bc44

Answer: A

NEW QUESTION 314
An administrator is using Amazon CloudFormation to deploy a three tier web application that consists of a web tier and application tier that will utilize Amazon DynamoDB for storage when creating the CloudFormation template which of the following would allow the application instance access to the DynamoDB tables without exposing API credentials?

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

Answer: C

NEW QUESTION 317
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 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 SAML2.0-compliam identity provider (IDP) to retrieve temporary security credentials to enable NOC members to sign in to the AWS Management Console.

Answer: D

NEW QUESTION 320
A benefits enrollment company is hosting a 3-tier web application running in a VPC on AWS which includes a NAT (Network Address Translation) instance in the public Web tier. There is enough provisioned capacity for the expected workload tor the new fiscal year benefit enrollment period plus some extra overhead Enrollment proceeds nicely for two days and then the web tier becomes unresponsive, upon investigation using CloudWatch and other monitoring tools it is discovered that there is an extremely large and unanticipated amount of inbound traffic coming from a set of 15 specific IP addresses over port 80 from a country where the benefits company has no customers. The web tier instances are so overloaded that benefit enrollment administrators cannot even SSH into them. Which actMty would be useful in defending against this attack?

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

Answer: D

NEW QUESTION 324
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Certified-Solutions-Architect-Professional Exam Questions With Answers.

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

https://www.2passeasy.com/dumps/AWS-Certified-Solutions-Architect-Professional/

Money Back Guarantee

AWS-Certified-Solutions-Architect-Professional Practice Exam Features:

* AWS-Certified-Solutions-Architect-Professional Questions and Answers Updated Frequently

* AWS-Certified-Solutions-Architect-Professional Practice Questions Verified by Expert Senior Certified Staff

* AWS-Certified-Solutions-Architect-Professional Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Certified-Solutions-Architect-Professional Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year