Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

NEW QUESTION 1
An organization is planning to host a Wordpress blog as well a joomla CMS on a single instance launched with VPC. The organization wants to have separate domains for each application and assign them using Route 53. The organization may have about ten instances each with two applications as mentioned above. While launching the instance, the organization configured two separate network interfaces (primary + ENI) and wanted to have two elastic IPs for that instance. It was suggested to use a public IP from AWS instead of an elastic IP as the number of elastic IPs is restricted. What action will you recommend to the organization?

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

Answer: B

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

NEW QUESTION 2
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 3
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 AutoScaIing policy for cost effectiveness.
Reference: http://media.amazonwebservices.com/AWS_Web_Hosting_Best_Practices.pdf

NEW QUESTION 4
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": "AIIow", "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"]
}| I
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 5
In which step of using AWS Direct Connect should the user determine the required port speed?

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

Answer: D

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

NEW QUESTION 6
In Amazon IAM, what is the maximum length for a role name?

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

Answer: C

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

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

A. ArnCheck
B. ArnMatch
C. ArnCase

D. ArnLike

Answer: D

Explanation:
Amazon Resource Name (ARN) condition operators let you construct Condition elements that restrict access based on comparing a key to an ARN. ArnLike, for instance, is a case-insensitive matching of the ARN. Each of the six colon-delimited components of the ARN is checked separately and each can include a multi-character match wildcard (*) or a single-character match wildcard (?).
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPoIicyLanguage_EIementDescriptions.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 15
What bandwidths do AWS Direct Connect currently support?

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

Answer: D

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

NEW QUESTION 17
The Principal element of an IAM policy refers to the specific entity that should be allowed or denied permission, whereas the translates to everyone except the specified entity.

A. NotPrincipa|
B. Vendor
C. Principal
D. Action

Answer: A

Explanation:
The element NotPrincipa| that is included within your IAM policy statements allows you to specify an exception to a list of principals to whom the access to a specific resource is either allowed or denied. Use the NotPrincipal element to specify an exception to a list of principals. For example, you can deny access to all principals except the one named in the NotPrincipa| element.

Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html#Principal

NEW QUESTION 19

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

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

Answer: A

Explanation:

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

NEW QUESTION 20

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

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

Answer: C

Explanation:

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

NEW QUESTION 25

An organization (account ID 123412341234) has configured the IAM policy to allow the user to modify his credentials. What will the below mentioned statement allow the user to perform?
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow", "Action": ["iam:AddUserToGroup",
"iam:RemoveUserFromGroup", "iam:GetGroup"
]!
"Resource": "arn:aws:iam:: 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": "2012-10-17",
"Statement": [{
"Effect": "Allow", "Action": ["iam:AddUserToGroup",
"iam:RemoveUserFromGroup", "iam:GetGroup"
]!
"Resource": "arn:aws:iam:: 123412341234:group/ TestingGroup "
}I
Reference:
http://docs.aws.amazon.com/IAM/latest/UserGuide/Credentials-Permissions-examples.html|#creds-polici es-credentials

NEW QUESTION 30

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 33
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:AssumeRole

Answer: D

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

NEW QUESTION 36
The MySecureData company has five branches across the globe. They want to expand their data centers such that their web server will be in the AWS and each branch would have their own database in the local data center. Based on the user login, the company wants to connect to the data center. How can MySecureData company implement this scenario with the AWS VPC?

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

Answer: B

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. If the organization has multiple VPN connections, he can provide secure communication between sites using the AWS VPN CIoudHub.
The VPN CIoudHub operates on a simple hub-and-spoke model that the user can use with or without a VPC. This design is suitable for customers with multiple branch offices and existing internet connections who would like to implement a convenient, potentially low-cost hub-and-spoke model for primary or backup connectMty between remote offices.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPN_CIoudHub.htmI

NEW QUESTION 38
How much memory does the cr1.8xIarge instance type provide?

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

Answer: D

Explanation:
The CR1 instances are part of the memory optimized instances. They offer lowest cost per GB RAM among all the AWS instance families. CR1 instances are part of the new generation of memory optimized instances, which can offer up to 244 GB RAM and run on faster CPUs (Intel Xeon E5-2670 with NUMA support) in comparison to the NI2 instances of the same family. They support cluster networking for bandwidth intensive applications. cr1.8x|arge is one of the largest instance types of the CR1 family, which can offer 244 GB RAM.
Reference: http://aws.amazon.com/ec2/instance-types/

NEW QUESTION 41
You want to define permissions for a role in an IAM policy. Which of the following configuration formats should you use?

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

Answer: C

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

NEW QUESTION 46
True or False: Amazon EIastiCache supports the Redis key-value store.

A. True, EIastiCache supports the Redis key-value store, but with limited functionalities.
B. False, EIastiCache 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 50
An organization is setting up an application on AWS to have both High Availabilty (HA) and Disaster Recovery (DR). The organization wants to have both Recovery point objective (RPO) and Recovery time objective (RTO) of 10 minutes. Which of the below mentioned service configurations does not help the organization achieve the said RPO and RTO?

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

Answer: C

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

NEW QUESTION 54
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 55
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/DeveloperGuide/what-is-datapipeline.html

NEW QUESTION 59
The CFO of a company wants to allow one of his employees to view only the AWS usage report page. Which of the below mentioned IAM policy statements allows the user to have access to the AWS usage report page?

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

Answer: C

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

NEW QUESTION 61
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 63
An organization has created 5 IAM users. The organization wants to give them the same login ID but different passwords. How can the organization achieve this?

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

Answer: C

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

NEW QUESTION 66
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 70
An organization is planning to setup a management network on the AWS VPC. The organization is trying to secure the webserver on a single VPC instance such that it allows the internet traffic as well as the back-end management traffic. The organization wants to make so that the back end management network interface can receive the SSH traffic only from a selected IP range, while the internet facing webserver will have an IP address which can receive traffic from all the internet IPs.
How can the organization achieve this by running web server on a single instance?

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

Answer: C

Explanation:

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

NEW QUESTION 73
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 75
ExamKiIler 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 empIoyees' 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 ExamKiI|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 78
An organization is undergoing a security audit. The auditor wants to view the AWS VPC configurations as the organization has hosted all the applications in the AWS VPC. The auditor is from a remote place and wants to have access to AWS to view all the VPC records.
How can the organization meet the expectations of the auditor without compromising on the security of their AWS infrastructure?

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

Answer: C

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

NEW QUESTION 82
Cognito Sync is an AWS service that you can use to synchronize user profile data across mobile devices without requiring your own backend. When the device is online, you can synchronize data. If you also set up push sync, what does it allow you to do?

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

Answer: C

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

NEW QUESTION 85
An organization is planning to create a secure scalable application with AWS VPC and ELB. The organization has two instances already running and each instance has an ENI attached to it in addition to a primary network interface. The primary network interface and additional ENI both have an elastic IP attached to it.
If those instances are registered with ELB and the organization wants ELB to send data to a particular EIP of the instance, how can they achieve this?

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

Answer: A

Explanation:
Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances. There are two ELBs available with VPC: internet facing and internal (private) ELB. For the internet facing ELB it is required that the ELB should be in a public subnet.
When the user registers a multi-homed instance (an instance that has an Elastic Network Interface (ENI) attached) with a load balancer, the load balancer will route the traffic to the IP address of the primary network interface (eth0).
Reference: http://docs.aws.amazon.com/E|asticLoadBalancing/latest/DeveloperGuide/gs-ec2VPC.html

NEW QUESTION 89
Which of the following cache engines does Amazon ElastiCache support?

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

Answer: A

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

NEW QUESTION 91
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 94
A bucket owner has allowed another account's IAM users to upload or access objects in his bucket. The IAM user of Account A is trying to access an object created by the IAM user of account B. What will happen in this scenario?

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

Answer: B

Explanation:
If a IAM user is trying to perform some action on an object belonging to another AWS user's bucket, S3 will verify whether the owner of the IAM user has given
sufficient permission to him. It also verifies the policy for the bucket as well as the policy defined by the object owner.

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

NEW QUESTION 97
An organization has hosted an application on the EC2 instances. There will be multiple users connecting to the instance for setup and configuration of application. The organization is planning to implement certain security best practices. Which of the below mentioned pointers will not help the organization achieve better security arrangement?

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

Answer: A

Explanation:
Since AWS is a public cloud any application hosted on EC2 is prone to hacker attacks. It becomes extremely important for a user to setup a proper security mechanism on the EC2 instances. A few of the security measures are listed below:
Always keep the OS updated with the latest patch
Always create separate users with in OS if they need to connect with the EC2 instances, create their keys and disable their password
Create a procedure using which the admin can revoke the access of the user when the business work on the EC2 instance is completed
Lock down unnecessary ports
Audit any proprietary applications that the user may be running on the EC2 instance
Provide temporary escalated prMleges, such as sudo for users who need to perform occasional prMleged tasks
The IAM is useful when users are required to work with AWS resources and actions, such as launching an instance. It is not useful to connect (RDP / SSH) with an instance.
Reference: http://aws.amazon.com/articles/1233/

NEW QUESTION 100
By default, temporary security credentials for an IAM user are valid for a maximum of 12 hours, but you can request a duration as long as hours.

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

Answer: B

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

NEW QUESTION 102
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 103
True or False : "|n the context of Amazon EIastiCache, from the application's point of view, connecting to the cluster configuration endpoint is no different than connecting directly to an indMdual cache node."

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

Answer: B

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

NEW QUESTION 106

An organization is setting up a highly scalable application using Elastic Beanstalk. They are using Elastic Load Balancing (ELB) as well as a Virtual Private Cloud (VPC) with public and private subnets. They have the following requirements:
. All the EC2 instances should have a private IP
. All the EC2 instances should receive data via the ELB's. Which of these will not be needed in this setup?

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

Answer: A

Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. If the organization wants the Amazon EC2 instances to have a private IP address, he should create a public and private subnet for VPC in each Availability Zone (this is an AWS Elastic Beanstalk requirement). The organization should add their public resources, such as ELB and NAT to the public subnet, and AWS Elastic Beanstalk will assign them unique elastic IP addresses (a static, public IP address). The organization should launch Amazon EC2 instances in a private subnet so that AWS Elastic Beanstalk assigns them non-routable private IP addresses. Now the organization should configure route tables with the following rules:
. route all inbound traffic from ELB to EC2 instances
. route all outbound traffic from EC2 instances through NAT
Reference: http://docs.aws.amazon.com/elasticbeanstaIk/latest/dg/AWSHowTo-vpc.html

NEW QUESTION 111
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 116
How does in-memory caching improve the performance of applications in E|astiCache?

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

Answer: D

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

NEW QUESTION 119
How can a user list the IAM Role configured as a part of the launch config?

A. as-describe-launch-configs --iam-profiIe
B. as-describe-launch-configs --show-Iong
C. as-describe-launch-configs —iam-role
D. as-describe-launch-configs —roIe

Answer: B

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

NEW QUESTION 124
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 128
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 129
Select the correct statement about Amazon EIastiCache.

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

Answer: A

Explanation:
EIastiCache is a web service that makes it easy to set up, manage, and scale a distributed in-memory cache environment in the cloud. It provides a high-performance, scalable, and cost-effective caching solution, while removing the complexity associated with deploying and managing a distributed cache environment. With EIastiCache, you can quickly deploy your cache environment, without having to provision hardware or install software.
Reference: http://docs.aws.amazon.com/AmazonE|astiCache/latest/UserGuide/WhatIs.html

NEW QUESTION 133
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/Iatest/UserGuide/VPC_Introduction.htm|

NEW QUESTION 137
What happens when Dedicated instances are launched into a VPC?

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

Answer: C

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

NEW QUESTION 140
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 144
You create a VPN connection, and your VPN device supports Border Gateway Protocol (BGP). Which of the following should be specified to configure the VPN connection?

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

Answer: C

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

NEW QUESTION 148
Can a Direct Connect link be connected directly to the Internet?

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

Answer: D

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

NEW QUESTION 152
ExamKiIIer has created a multi-tenant Learning Management System (LMS). The application is hosted for five different tenants (clients) in the VPCs of the respective AWS accounts of the tenant. ExamKiIIer wants to setup a centralized server which can connect with the LMS of each tenant upgrade if required. ExamKiIIer also wants to ensure that one tenant VPC should not be able to connect to the other tenant VPC for security reasons. How can ExamKi|er setup this scenario?

A. ExamKiI|er has to setup one centralized VPC which will peer in to all the other VPCs of the tenants.
B. ExamKiI|er should setup VPC peering with all the VPCs peering each other but block the IPs from CIDR of the tenant VPCs to deny them.
C. ExamKiI|er should setup all the VPCs with the same CIDR but have a centralized VP
D. This way only the centralized VPC can talk to the other VPCs using VPC peering.
E. ExamKiI|er should setup all the VPCs meshed together with VPC peering for all 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. The organization wants to setup that one VPC can connect with all the other VPCs but all other VPCs cannot connect among each other. This can be achieved by configuring VPC peering where one VPC is peered with all the other VPCs, but the other VPCs are not peered to each other. The VPCs are in the same or a separate AWS account and should not have overlapping CIDR blocks.
Reference:
http://docs.aws.amazon.com/AmazonVPC/Iatest/PeeringGuide/peering-configurations-full-access.html# many-vpcs-full-acces

NEW QUESTION 157
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 160
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 165
In the context of IAM roles for Amazon EC2, which of the following NOT true about delegating permission to make API requests?

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

Answer: A

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

NEW QUESTION 168
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 169
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 172
How does AWS Data Pipeline execute actMties on on-premise resources or AWS resources that you manage?

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

Answer: A

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

NEW QUESTION 173
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/Whatls.html

NEW QUESTION 177
What is a possible reason you would need to edit claims issued in a SAML token?

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

Answer: A

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

NEW QUESTION 179
A government client needs you to set up secure cryptographic key storage for some of their extremely confidential data. You decide that the AWS CIoudHSM is the best service for this. However, there seem to be a few pre-requisites before this can happen, one of those being a security group that has certain ports open. Which of the following is correct in regards to those security groups?

A. A security group that has 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 C|oudHSM requires the following environment before an HSM appliance can be provisioned. A virtual private cloud (VPC) in the region where you want the AWS 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 CIoudHSM.
An EC2 instance, in the same VPC as the HSM appliance, that has the SafeNet client software installed. This instance is referred to as the control instance and is used to connect to and manage the HSM appliance.
A security group that has port 22 (for SSH) or port 3389 (for RDP) open to your network. This security group is attached to your control instances so you can access them remotely.

NEW QUESTION 180
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 CIoudFront

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/AmazonCloudFront/latest/DeveloperGuide/Troubleshooting.html

NEW QUESTION 181
Mike is appointed as Cloud Consultant in ExamKiller.com. ExamKiller has the following VPCs set-up in the US East Region:
A VPC with CIDR block 10.10.0.0/16, a subnet in that VPC with CIDR block 10.10.1.0/24 A VPC with CIDR block 10.40.0.0/16, a subnet in that VPC with CIDR block 10.40.1.0/24
ExamKiller.com is trying to establish network connection between two subnets, a subnet with CIDR block 10.10.1.0/24 and another subnet with CIDR block 10.40.1.0/24. Which one of the following solutions should Mike recommend to ExamKiller.com?

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

Answer: C

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

NEW QUESTION 184
A user is hosting a public website on AWS. The user wants to have the database and the app server on the AWS VPC. The user wants to setup a database that can connect to the Internet for any patch upgrade but cannot receive any request from the internet. How can the user set this up?

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

Answer: D

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

NEW QUESTION 189
In Amazon ElastiCache, the default cache port is:

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

Answer: D

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

NEW QUESTION 194
A user has created a VPC with public and private subnets using the VPC wizard. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.0.0/24 . The NAT instance ID is i-a12345. Which of the below mentioned entries are required in the main route table attached with the private subnet to allow instances to connect with the internet?

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

Answer: D

Explanation:
A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the Internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create two route tables and attach to the subnets. The main route table will have the entry "Destination: 0.0.0.0/0 and Target: i-a12345", which allows all the instances in the private subnet to connect to the internet using NAT.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

NEW QUESTION 196
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 199
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 204
Identify a true statement about the statement ID (Sid) in IAM.

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

Answer: A

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

NEW QUESTION 206
In Amazon ElastiCache, which of the following statements is correct?

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

Answer: B

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

NEW QUESTION 207
An organization has setup RDS with VPC. The organization wants RDS to be accessible from the internet. Which of the below mentioned configurations is not required in this scenario?

A. The organization must enable the parameter in the console which makes the RDS instance publicly accessible.
B. The organization must allow access from the internet in the RDS VPC security group,
C. The organization must setup RDS with the subnet group which has an external IP.
D. The organization must enable the VPC attributes DNS hostnames and DNS resolutio

Answer: C

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

NEW QUESTION 212
An organization, which has the AWS account ID as Q99988887777, has created 50 IAM users. All the users are added to the same group examkiller. If the organization has enabled that each IAM user can login with the AWS console, which AWS login URL will the IAM users use??

A. https://Q99988887777.aws.amazon.com/examkiIler/
B. https://signin.aws.amazon.com/examki|Ier/
C. https://examkiller.signin.aws.amazon.com/999988887777/console/
D. https://999988887777.signin.aws.amazon.com/console/

Answer: D

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

NEW QUESTION 213
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 217
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 CIoudWatch and RDS with read replicas.
C. Stateful instances for the web and application tier in an autoscaling group monitored with CIoudWatc
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 222
You are designing a photo-sharing mobile app. The application will store all pictures in a single Amazon S3 bucket. Users will upload pictures from their mobile device directly to Amazon S3 and will be able to view and download their own pictures directly from Amazon S3.
You want to configure security to handle potentially millions of users in the most secure manner possible. What should your server-side application do when a new user registers on the photo-sharing mobile application?

A. Create an IAM use
B. Update the bucket policy with appropriate permissions for the IAM use
C. Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.
D. Create an IAM use

E. Assign appropriate permissions to the IAM use

F. Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.

G. Create a set of long-term credentials using AWS Security Token Service with appropriate permission

H. Store these credentials in the mobile app and use them to access Amazon S3.

I. Record the user's information in Amazon RDS and create a role in IAM with appropriate permission

J. When the user uses their mobile app, create temporary credentials using the AWS Security Token Service "AssumeRole" functio

K. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.

L. Record the user's information in Amazon DynamoD

M. When the user uses their mobile app, create temporary credentials using AWS Security Token Service with appropriate permission

N. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.

Answer: D

NEW QUESTION 223
You have a periodic Image analysis application that gets some files In Input analyzes them and tor each file writes some data in output to a ten file the number of files in input per day is high and concentrated in a few hours of the day.
Currently you have a server on EC2 with a large EBS volume that hosts the input data and the results it takes almost 20 hours per day to complete the process What services could be used to reduce the elaboration time and improve the availability of the solution?

A. S3 to store I/O file
B. SQS to distribute elaboration commands to a group of hosts working in paralle
C. Auto scaling to dynamically size the group of hosts depending on the length of the SQS queue
D. EBS with Provisioned IOPS (PIOPS) to store I/O file
E. SNS to distribute elaboration commands to a group of hosts working in parallel Auto Scaling to dynamically size the group of hosts depending on the number of SNS notifications
F. S3 to store I/O files, SNS to distribute evaporation commands to a group of hosts working in paralle
G. Auto scaling to dynamically size the group of hosts depending on the number of SNS notifications
H. EBS with Provisioned IOPS (PIOPS) to store I/O files SQS to distribute elaboration commands to a group of hosts working in parallel Auto Scaling to dynamically size the group ot hosts depending on the length of the SQS queue.

Answer: D

NEW QUESTION 224
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 c3.8xIarge 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
You are the new IT architect in a company that operates a mobile sleep tracking application.
When activated at night, the mobile app is sending collected data points of 1 kilobyte every 5 minutes to your backend.
The backend takes care of authenticating the user and writing the data points into an Amazon DynamoDB table.
Every morning, you scan the table to extract and aggregate last night's data on a per user basis, and store the results in Amazon S3. Users are notified via Amazon SNS mobile push notifications that new data is available, which is parsed and visualized by the mobile app.
Currently you have around 100k users who are mostly based out of North America. You have been tasked to optimize the architecture of the backend system to lower cost. What would you recommend? Choose 2 answers

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

Answer: AD

NEW QUESTION 231
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

Answer: A

NEW QUESTION 244
You would like to create a mirror image of your production environment in another region for disaster recovery purposes. Which of the following AWS resources do not need to be recreated in the second region? (Choose 2 answers)

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

Answer: AC

NEW QUESTION 248
A customer has a 10 GB AWS Direct Connect connection to an AWS region where they have a web application hosted on Amazon Elastic Computer Cloud (EC2). The application has dependencies on an on-premises mainframe database that uses a BASE (Basic Available. Sort stale Eventual consistency) rather than an ACID (Atomicity. Consistency isolation. Durability) consistency model. The application is exhibiting undesirable behavior because the database is not able to handle the volume of writes. How can you reduce the load on your on-premises database resources in the most cost-effective way?

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

Answer: A

NEW QUESTION 250
You are responsible for a legacy web application whose server environment is approaching end of life You would like to migrate this application to AWS as quickly as possible, since the application environment currently has the following limitations:
The VM's single 10GB VNIDK is almost full; Nle virtual network interface still uses the 10IV|bps driver, which leaves your 100Mbps WAN connection completely underutilized;
It is currently running on a highly customized. Windows VM within a VMware environment; You do not have me installation media;
This is a mission critical application with an RTO (Recovery Time Objective) of 8 hours. RPO (Recovery Point Objective) of 1 hour. How could you best migrate this application to AWS while meeting your business continuity requirements?

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

Answer: A

NEW QUESTION 251
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 253
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 256

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

A. Use AWS data Pipeline to schedule a DynamoDB cross region copy once a 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 258
You must architect the migration of a web application to AWS. The application consists of Linux web servers running a custom web server. You are required to save the logs generated from the application to a durable location.
What options could you select to migrate the application to AWS? (Choose 2)

A. Create an AWS Elastic Beanstalk application using the custom web server platfor
B. Specify the web server executable and the application project and source file
C. Enable log file rotation to Amazon Simple Storage Service (S3).
D. Create Dockerfile for the applicatio
E. Create an AWS OpsWorks stack consisting of a custom laye
F. Create custom recipes to install Docker and to deploy your Docker container using the Dockerfil
G. Create customer recipes to install and configure the application to publish the logs to Amazon CloudWatch Logs.
H. Create Dockerfile for the applicatio
I. Create an AWS OpsWorks stack consisting of a Docker layer that uses the Dockerfil
J. Create custom recipes to install and configure Amazon Kineses to publish the logs into Amazon CloudWatch.
K. Create a Dockerfile for the applicatio
L. Create an AWS Elastic Beanstalk application using the Docker platform and the Dockerfil
M. Enable logging the Docker configuration to automatically publish the application log
N. Enable log file rotation to Amazon S3.
O. Use VM import/Export to import a virtual machine image of the server into AWS as an AM
P. Create an Amazon Elastic Compute Cloud (EC2) instance from AMI, and install and configure the Amazon CloudWatch Logs agen
Q. Create a new AMI from the instanc
R. Create an AWS Elastic Beanstalk application using the AMI platform and the new AMI.

Answer: AD

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 264
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 267

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 268

You are running a news website in the eu-west-1 region that updates every 15 minutes. The website has a world-wide audience it uses an Auto Scaling group behind an Elastic Load Balancer and an Amazon
RDS database Static content resides on Amazon S3, and is distributed through Amazon CloudFront. Your Auto Scaling group is set to trigger a scale up event at 60% CPU utilization, you use an Amazon RDSextra large DB instance with 10.000 Provisioned IOPS its CPU utilization is around 80%. While freeable memory is in the 2 GB range.

Web analytics reports show that the average load time of your web pages is around 1.5 to 2 seconds, but your SEO consultant wants to bring down the average load time to under 0.5 seconds.

How would you improve page load times for your users? (Choose 3 answers)

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

Answer: ABD

NEW QUESTION 269

A corporate web application is deployed within an Amazon Virtual Private Cloud (VPC) and is connected to the corporate data center via an IPSec VPN. The application must authenticate against the
on-premises LDAP server. After authentication, each logged-in user can only access an Amazon Simple Storage Space (S3) keyspace specific to that user.
Which two approaches can satisfy these objectives? (Choose 2 answers)

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

Answer: BC

NEW QUESTION 273

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 276

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 277

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 281

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

You are designing a personal document-archMng solution for your global enterprise with thousands of employee. Each employee has potentially gigabytes of data to be backed up in this archMng solution. The solution will be exposed to the employees as an application, where they can just drag and drop their files to the archMng system. Employees can retrieve their archives through a web interface. The corporate network has high bandwidth AWS Direct Connect connectMty to AWS.

You have a regulatory requirement that all data needs to be encrypted before being uploaded to the cloud.

How do you implement this in a highly available and cost-efficient way?

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

Answer: C

NEW QUESTION 290

You are designing a connectMty solution between on-premises infrastructure and Amazon VPC. Your servers on-premises will be communicating with your VPC instances. You will be establishing IPSec tunnels over the Internet You will be using VPN gateways, and terminating the IPSec tunnels on AWS supported customer gateways.

Which of the following objectives would you achieve by implementing an IPSec tunnel as outlined above? Choose 4 answers

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

Answer: CDEF

NEW QUESTION 291

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

What should you do in order to avoid errors for future deployments? (Choose 2 answer)

A. Add an Elastic Load Balancing health check to the Auto Scaling grou
B. Set a short period for the health checks to operate as soon as possible in order to prevent premature registration of the instance to theload balancer.

C. Enable EC2 instance C|oudWatch alerts to change the launch configuration's AMI to the previous on
D. Gradually terminate instances that are using the new AMI.
E. Set the Elastic Load Balancing health check configuration to target a part of the application that fully tests application health and returns an error if the tests fail.
F. Create a new launch configuration that refers to the new AMI, and associate it with the grou
G. Double the size of the group, wait for the new instances to become healthy, and reduce back to the original size.If new instances do not become healthy, associate the previous launch configuration.
H. Increase the Elastic Load Balancing Unhealthy Threshold to a higher value to prevent an unhealthy instance from going into service behind the load balancer.

Answer: CD

NEW QUESTION 296
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

AWS-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

100% Actual & Verified — Instant Download, Please Click
Order The AWS-Certified-Solutions-Architect-Professional Practice Test Here