Exam Questions AWS-Certified-Solutions-Architect-Professional

AWS-Certified-Solutions-Architect-Professional

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

NEW QUESTION 1

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

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

Answer: D

Explanation: By default, Amazon Cognito maintains the last-written version of the data. You can override this behavior and resolve data conflicts programmatically. In addition, push synchronization allows you to use Amazon Cognito to send a silent push notification to all devices associated with an identity to notify them that new data is available.
Reference: http://aws.amazon.com/cognito/faqs/

NEW QUESTION 2

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

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

Answer: C

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

NEW QUESTION 3

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

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

Answer: B

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

NEW QUESTION 4

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

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

Answer: B

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

NEW QUESTION 5

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 6
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 7
A customer has a website which shows all the deals available across the market. The site experiences a load of 5 large EC2 instances generally. However, a week before ThanksgMng vacation they encounter a load of almost 20 large instances. The load during that period varies over the day based on the office timings. Which of the below mentioned solutions is cost effective as well as help the website achieve better performance?

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

Answer: D

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

NEW QUESTION 8
An organization is setting a website on the AWS VPC. The organization has blocked a few IPs to avoid a D-DOS attack. How can the organization configure that a request from the above mentioned IPs does not access the application instances?

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

Answer: D

Explanation: A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. AWS provides two features that the user can use to increase security in VPC: security groups and network ACLs. Security group works at the instance level while ACL works at the subnet level. ACL allows both allow and deny rules.
Thus, when the user wants to reject traffic from the selected IPs it is recommended to use ACL with subnets.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_ACLs.html

NEW QUESTION 9
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 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 10
An organization is planning to host an application on the AWS VPC. The organization wants dedicated instances. However, an AWS consultant advised the organization not to use dedicated instances with VPC as the design has a few limitations. Which of the below mentioned statements is not a limitation of dedicated instances with VPC?

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

Answer: C

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

NEW QUESTION 10
You have subscribed to the AWS Business and Enterprise support plan. Your business has a backlog of problems, and you need about 20 of your IAM users to open technical support cases. How many users can open technical support cases under the AWS Business and Enterprise support plan?

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

Answer: C

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

NEW QUESTION 14
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 19
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 21
An organization is creating a VPC for their application hosting. The organization has created two private subnets in the same AZ and created one subnet in a separate zone. The organization wants to make a
HA system with the internal ELB. Which of these statements is true with respect to an internal ELB in this scenario?

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

Answer: A

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

NEW QUESTION 25
In Amazon EIastiCache, the failure of a single cache node can have an impact on the availability of your application and the load on your back-end database while EIastiCache provisions a replacement for the failed cache node and it get repopulated. Which of the following is a solution to reduce this potential availability impact?

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

Answer: B

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

NEW QUESTION 30
lV|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 34
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 36
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 39
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_eIements.htmI#PrincipaI

NEW QUESTION 42
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 46
A user is configuring MySQL RDS with PIOPS. What should be the minimum size of DB storage provided by the user?

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

Answer: D

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

NEW QUESTION 51
The Statement element, of an AWS IAM policy, contains an array of indMdual statements. Each indMdual statement is a(n) block enclosed in braces { }.

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

Answer: C

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

NEW QUESTION 54
If no explicit deny is found while applying IAM's Policy Evaluation Logic, the enforcement code looks for any instructions that would apply to the request.

A. "cancel"
B. "suspend"
C. "a||ow"
D. "valid"

Answer: C

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

NEW QUESTION 58
An organization (account ID 123412341234) has configured the IAM policy to allow the user to modify his credentials. What will the below mentioned statement allow the user to perform?
{
"Version": "2012-10-I7",
"Statement": [{
"Effect": "A||ow", "Action": ["iam:AddUserToGroup",
"iam:RemoveUserFromGroup", "iam:GetGroup"
]!
"Resource": "arn:aws:iam:: I23412341234:group/TestingGroup"
}I

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

Answer: A

Explanation: AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the organization (account ID 123412341234) wants their users to manage their subscription to the groups, they should create a relevant policy for that. The below mentioned policy allows the respective IAM user to update the membership of the group called MarketingGroup.
{
"Version": "2012-10-I7",
"Statement": [{
"Effect": "A||ow", "Action": ["iam:AddUserToGroup",
"iam:RemoveUserFromGroup", "iam:GetGroup"
]!
"Resource": "arn:aws:iam:: 123412341234:group/ TestingGroup "
}I
Reference:
http://docs.aws.amazon.com/IAM/latest/UserGuide/CredentiaIs-Permissions-examples.htm|#creds-po|ici es-credentials

NEW QUESTION 61
One of your AWS Data Pipeline actMties has failed consequently and has entered a hard failure state after retrying thrice. You want to try it again. Is it possible to increase the number of automatic retries to more than thrice?

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

Answer: D

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

NEW QUESTION 65
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 68
How many cg1.4xlarge on-demand instances can a user run in one region without taking any limit increase approval from AWS?

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

Answer: B

Explanation: Generally AWS EC2 allows running 20 on-demand instances and 100 spot instances at a time. This limit can be increased by requesting at https://aws.amazon.com/contact-us/ec2-request. Excluding certain types of instances, the limit is lower than mentioned above. For cg1.4xlarge, the user can run only 2
on-demand instances at a time.
Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#|imits_ec2

NEW QUESTION 72
Regarding Amazon SNS, you can send notification messages to mobile devices through any of the following supported push notification services, EXCEPT:

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

Answer: A

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

NEW QUESTION 76
IAM Secure And Scalable is an organization which provides scalable and secure SAAS to its clients. They are planning to host a web server and App server on AWS VPC as separate tiers. The organization wants to implement the scalability by configuring Auto Scaling and load balancer with their app servers (middle tier) too. Which of the below mentioned options suits their requirements?

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

Answer: D

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

NEW QUESTION 79
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 80
Does Amazon RDS API provide actions to modify DB instances inside a VPC and associate them with DB Security Groups?

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

Answer: B

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

NEW QUESTION 82
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.cIoudfront.net/AWS_Disaster_Recovery.pdf

NEW QUESTION 84
By default, what is the maximum number of Cache Nodes you can run in Amazon EIastiCache?

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

Answer: A

Explanation: In Amazon EIastiCache, you can run a maximum of 20 Cache Nodes. Reference: http://aws.amazon.com/e|asticache/faqs/

NEW QUESTION 85
Does an AWS Direct Connect location provide access to Amazon Web Services in the region it is associated with as well as access to other US regions?

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

Answer: C

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

NEW QUESTION 86
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 90
What feature of the load balancing service attempts to force subsequent connections to a service to be redirected to the same node as long as it is online?

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

Answer: D

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

NEW QUESTION 95
What types of identities do Amazon Cognito identity pools support?

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

Answer: A

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

NEW QUESTION 100
In IAM, which of the following is true of temporary security credentials?

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

Answer: A

Explanation: Temporary credentials in IAM are valid throughout their defined duration of time and hence can't be revoked. However, because permissions are evaluated each time an AWS request is made using the credentials, you can achieve the effect of revoking the credentials by changing the permissions for the credentials even after they have been issued. Reference:
http://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp_control-access_disable-perms.h tml

NEW QUESTION 101
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 105
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.html

NEW QUESTION 108
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 113
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.html

NEW QUESTION 118
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 120
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 121
What is the maximum length for an instance profile name in AWS IAM?

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

Answer: B

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

NEW QUESTION 122
In Amazon Cognito, your mobile app authenticates with the Identity Provider (|dP) using the provider's SDK. Once the end user is authenticated with the IdP, the OAuth or OpenID Connect token returned from the IdP is passed by your app to Amazon Cognito, which returns a new for the user and a set of temporary, limited-prMlege AWS credentials.

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

Answer: C

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

NEW QUESTION 124
What is the maximum length for a certificate ID in AWS IAM?

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

Answer: D

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

NEW QUESTION 125
A user has configured EBS volume with PIOPS. The user is not experiencing the optimal throughput. Which of the following could not be factor affecting I/O performance of that EBS volume?

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

Answer: C

Explanation: If the user is not experiencing the expected IOPS or throughput that is provisioned, ensure that the EC2 bandwidth is not the limiting factor, the instance is EBS-optimized (or include 10 Gigabit network
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 129
If a single condition within an IAM policy includes multiple values for one key, it will be evaluated using a logical .

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

Answer: A

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

NEW QUESTION 131
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 133
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 135
In a VPC, can you modify a set of DHCP options after you create them?

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

Answer: C

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

NEW QUESTION 139
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.html

NEW QUESTION 141
Which statement is NOT true about a stack which has been created in a Virtual Private Cloud (VPC) in AWS OpsWorks?

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

Answer: A

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

NEW QUESTION 145
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 150
One of the AWS account owners faced a major challenge in June as his account was hacked and the hacker deleted all the data from his AWS account. This resulted in a major blow to the business.
Which of the below mentioned steps would not have helped in preventing this action?

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

Answer: C

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

NEW QUESTION 154
With Amazon Elastic MapReduce (Amazon EMR) you can analyze and process vast amounts of data. The cluster is managed using an open-source framework called Hadoop.
You have set up an application to run Hadoop jobs. The application reads data from DynamoDB and generates a temporary file of 100 TBs.
The whole process runs for 30 minutes and the output of the job is stored to S3. Which of the below mentioned options is the most cost effective solution in this case?

A. Use Spot Instances to run Hadoop jobs and configure them with EBS volumes for persistent data storage.

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

Answer: B

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

NEW QUESTION 157
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 161
True or False : "|n the context of Amazon ElastiCache, from the application's point of view, connecting to the cluster configuration endpoint is no different than connecting directly to an 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 165
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 167
An EC2 instance that performs source/destination checks by default is launched in a private VPC subnet. All security, NACL, and routing definitions are configured as expected. A custom NAT instance is launched.
Which of the following must be done for the custom NAT instance to work?

A. The source/destination checks should be disabled on the NAT instance.
B. The NAT instance should be launched in public subnet.

C. The NAT instance should be configured with a public IP address.
D. The NAT instance should be configured with an elastic IP addres

Answer: A

Explanation: Each EC2 instance performs source/destination checks by default. This means that the instance must be the source or destination of any traffic it sends or receives. However, a NAT instance must be able to send and receive traffic when the source or destination is not itself. Therefore, you must disable source/destination checks on the NAT instance.
Reference:
http://docs.aws.amazon.com/AmazonVPC/Iatest/UserGuide/VPC_NAT_|nstance.htm|#EIP_Disab|e_Src DestCheck

NEW QUESTION 172
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/NlultiplelP.htmI

NEW QUESTION 173
An organization is making software for the CIA in US

A. CIA agreed to host the application on AWS but ina secure environmen
B. The organization is thinking of hosting the application on the AWS GovC|oud regio
C. Which of the below mentioned difference is not correct when the organization is hosting on the AWS GovCIoud in comparison with the AWS standard region?
D. The billing for the AWS GovCLoud will be in a different account than the Standard AWS account.
E. GovCIoud region authentication is isolated from Amazon.com.
F. Physical and logical administrative access only to U.
G. persons.
H. It is physically isolated and has logical network isolation from all the other region

Answer: A

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

NEW QUESTION 178
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 EIastiCache, 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 183
A user is thinking to use EBS PIOPS volume. Which of the below mentioned options is a right use case for the PIOPS EBS volume?

A. Analytics
B. System boot volume
C. Nlongo DB

D. Log processing

Answer: C

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

NEW QUESTION 188
Which of the following is true of an instance profile when an IAM role is created using the console?

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

Answer: B

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

NEW QUESTION 189
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 191
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 196
Attempts, one of the three types of items associated with the schedule pipeline in the AWS Data Pipeline, provides robust data management.
Which of the following statements is NOT true about Attempts?

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

Answer: C

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

NEW QUESTION 199

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

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

Answer: C

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

NEW QUESTION 203
AWS Direct Connect itself has NO specific resources for you to control access to. Therefore, there are no AWS Direct Connect Amazon Resource Names (ARNs) for you to use in an Identity and Access Nlanagement (IAM) policy. With that in mind, how is it possible to write a policy to control access to AWS Direct Connect actions?

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

Answer: C

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

NEW QUESTION 205
Identify an application that polls AWS Data Pipeline for tasks and then performs those tasks.

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

Answer: C

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

NEW QUESTION 210
With respect to AWS Lambda permissions model, at the time you create a Lambda function, you specify an IAM role that AWS Lambda can assume to execute your Lambda function on your behalf. This role is also referred to as the role.

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

Answer: B

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

NEW QUESTION 215
Within an IAM policy, can you add an IfExists condition at the end of a Null condition?

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

Answer: C

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

NEW QUESTION 220
Regarding Identity and Access Management (IAM), Which type of special account belonging to your application allows your code to access Google services programmatically?

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

Answer: A

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

NEW QUESTION 222
IAM users do not have permission to create Temporary Security Credentials for federated users and roles by default. In contrast, IAM users can call without the need of any special permissions

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

Answer: C

Explanation: Currently the STS API command GetSessionToken is available to every IAM user in your account without previous permission. In contrast, the GetFederationToken command is restricted and explicit permissions need to be granted so a user can issue calls to this particular Action
Reference: http://docs.aws.amazon.com/STS/latest/UsingSTS/STSPermission.html

NEW QUESTION 227
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 232
An organization has developed an application which provides a smarter shopping experience. They need to show a demonstration to various stakeholders who may not be able to access the in premise
application so they decide to host a demo version of the application on AWS. Consequently they will need a fixed elastic IP attached automatically to the instance when it is launched.
In this scenario which of the below mentioned options will not help assign the elastic IP automatically?

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

Answer: A

Explanation: EC2 allows the user to launch On-Demand instances. If the organization is using an application temporarily only for demo purposes the best way to assign an elastic IP would be:

Launch an instance with a VPC and assign an EIP to the primary network interface. This way on every instance start it will have the same IP Create a bootstrapping script and provide it some metadata, such as user data which can be used to assign an EIP Create a controller instance which can schedule the start and stop of the instance and provide an EIP as a parameter so that the controller instance can check the instance boot and assign an EIP The instance metadata gives the current instance data, such as the public/private IP. It can be of no use for assigning an EIP.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AESDG-chapter-instancedata.html

NEW QUESTION 233
An organization is having a VPC for the HR department, and another VPC for the Admin department. The HR department requires access to all the instances running in the Admin VPC while the Admin department requires access to all the resources in the HR department. How can the organization setup this scenario?

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

Answer: A

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

NEW QUESTION 235
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 ExamKiIIer 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.htmI# many-vpcs-full-acces

NEW QUESTION 237
A user has created a VPC with CIDR 20.0.0.0/16. The user has created one subnet with CIDR 20.0.0.0/16 in this VPC. The user is trying to create another subnet with the same VPC for CIDR 20.0.0.1/24. What will happen in this scenario?

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

Answer: C

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

NEW QUESTION 239
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 241
In Amazon Cognito what is a silent push notification?

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

Answer: A

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

NEW QUESTION 244
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 "NumericLessThanEquaIs"?

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

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

NEW QUESTION 245
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 248
Which of following IAM policy elements lets you specify an exception to a list of actions?

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

Answer: D

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

NEW QUESTION 250
A user is configuring MySQL RDS with PIOPS. What should be the minimum PIOPS that the user should provision?

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

Answer: A

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

NEW QUESTION 251
You are setting up some EBS volumes for a customer who has requested a setup which includes a RAID (redundant array of inexpensive disks). AWS has some recommendations for RAID setups. Which RAID setup is not recommended for Amazon EBS?

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

Answer: C

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

NEW QUESTION 254
Once the user has set ElastiCache for an application and it is up and running, which services, does Amazon not provide for the user:

A. The ability for client programs to automatically identify all of the nodes in a cache cluster, and to initiate and maintain connections to all of these nodes
B. Automating common administrative tasks such as failure detection and recovery, and software patching
C. Providing default Time To Live (TTL) in the AWS Elasticache Redis Implementation for different type of data.
D. Providing detailed monitoring metrics associated with your Cache Nodes, enabling you to diagnose and react to issues very quickly

Answer: C

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

NEW QUESTION 259
True or False: In Amazon ElastiCache, you can use Cache Security Groups to configure the cache clusters that are part of a VPC.

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

Answer: A

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

NEW QUESTION 261
What is the role of the PollForTask action when it is called by a task runner in AWS Data Pipeline?

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

Answer: C

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

NEW QUESTION 265
What is the average queue length recommended by AWS to achieve a lower latency for the 200 PIOPS EBS volume?

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

Answer: B

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

NEW QUESTION 266
An organization is planning to host a web application in the AWS VPC. The organization does not want to host a database in the public cloud due to statutory requirements. How can the organization setup in this scenario?

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

Answer: A

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

NEW QUESTION 270
A user is planning to host a Highly Available system on the AWS VPC. Which of the below mentioned statements is helpful in this scenario?

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

Answer: A

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

NEW QUESTION 274
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-samI-claims-customization/

NEW QUESTION 276
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 CloudHSM.
An EC2 instance, in the same VPC as the HSM appliance, that has the SafeNet client software installed. This instance is referred to as the control instance and is used to connect to and manage the HSM appliance.
A security group that has port 22 (for SSH) or port 3389 (for RDP) open to your network. This security group is attached to your control instances so you can access them remotely.

NEW QUESTION 279
What is the network performance offered by the c4.8xlarge instance in Amazon EC2?

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

Answer: D

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

NEW QUESTION 284
An organization is setting up a web application with the JEE stack. The application uses the JBoss app server and |V|ySQL DB. The application has a logging module which logs all the actMties whenever a business function of the JEE application is called. The logging actMty takes some time due to the large size of the log file. If the application wants to setup a scalable infrastructure which of the below mentioned options will help achieve this setup?

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

Answer: D

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

NEW QUESTION 285
You're trying to delete an SSL certificate from the IAM certificate store, and you're getting the message "Certificate: <certificate-id> is being used by CloudFront."
Which of the following statements is probably the reason why you are getting this error?

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

Answer: C

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

NEW QUESTION 288
A user has set the IAM policy where it denies all requests if a request is not from IP 10.10.10.1/32. The other policy says allow all requests between 5 PM to 7 PM.
What will happen when a user is requesting access from IP 55.109.10.12/32 at 6 PM?

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

Answer: A

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

In this case since there are explicit deny and explicit allow statements. Thus, the request will be denied since deny overrides allow.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_EvaluationLogic.html

NEW QUESTION 291
Which of the following AWS services can be used to define alarms to trigger on a certain actMty, such as actMty success, failure, or delay in AWS Data Pipeline?

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

Answer: C

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

NEW QUESTION 296
You want to use Amazon Redshift and you are planning to deploy dw1.8xlarge nodes. What is the minimum amount of nodes that you need to deploy with this kind of configuration?

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

Answer: D

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

NEW QUESTION 300
Mike is appointed as Cloud Consultant in ExamKi|Ier.com. ExamKiI|er 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
ExamKiIIer.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 IV|ike recommend to ExamKiI|er.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 304
Can Provisioned IOPS be used on RDS instances launched in a VPC?

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

Answer: B

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

NEW QUESTION 309
To get started using AWS Direct Connect, in which of the following steps do you configure Border Gateway Protocol (BGP)?

A. Complete the Cross Connect

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

Answer: C

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

NEW QUESTION 313
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 315
An organization is setting up their website on AWS. The organization is working on various security measures to be performed on the AWS EC2 instances. Which of the below mentioned security mechanisms will not help the organization to avoid future data leaks and identify security weaknesses?

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

Answer: C

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

NEW QUESTION 317
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.htm|

NEW QUESTION 321
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 325
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

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

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

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

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

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