Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

NEW QUESTION 1
An IAM user is trying to perform an action on an object belonging to some other root account's bucket. Which of the below mentioned options will AWS S3 not verify?

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

Answer: B

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

NEW QUESTION 2
In the context of AWS IAM, identify a true statement about user passwords (login profiles).

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

Answer: B

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

NEW QUESTION 3
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 4
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 5
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 6
In which step of using AWS Direct Connect should the user determine the required port speed?

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

Answer: D

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

NEW QUESTION 7
While implementing the policy keys in AWS Direct Connect, if you use and the request comes from
an Amazon EC2 instance, the instance's public IP address is evaluated to determine if access is allowed.

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

Answer: C

Explanation:
While implementing the policy keys in Amazon RDS, if you use aws:SourceIp and the request comes from an Amazon EC2 instance, the instance's public IP address is evaluated to determine if access is allowed. Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/using_iam.html

NEW QUESTION 8
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/AccessPolicyLanguage_ElementDescriptions.html

NEW QUESTION 9
In Amazon ElastiCache, the failure of a single cache node can have an impact on the availability of your application and the load on your back-end database while ElastiCache provisions a replacement for the failed cache node and it get repopulated. Which of the following is a solution to reduce this potential availability impact?

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

Answer: B

Explanation:
In Amazon ElastiCache, the number of cache nodes in the cluster is a key factor in the availability of your cluster running Memcached. The failure of a single cache node can have an impact on the availability of your application and the load on your back-end database while ElastiCache provisions a replacement for the failed cache node and it get repopulated. You can reduce this potential availability impact by spreading your memory and compute capacity over a larger number

of cache nodes, each with smaller capacity, rather than using a fewer number of high capacity nodes.
Reference: http://docs.aws.amazon.com/AmazonEIastiCache/Iatest/UserGuide/CacheNode.Memcached.html

NEW QUESTION 10
IV|apMySite is setting up a web application in the AWS VPC. The organization has decided to use an AWS RDS instead of using its own DB instance for HA and DR requirements.
The organization also wants to secure RDS access. How should the web application be setup with RDS?

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

Answer: B

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

NEW QUESTION 10
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 11
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 13
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 16
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 19
An organization is hosting a scalable web application using AWS. The organization has configured ELB and Auto Scaling to make the application scalable. Which of the below mentioned statements is not required to be followed for ELB when the application is planning to host a web application on VPC?

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

Answer: A

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

NEW QUESTION 21
How can multiple compute resources be used on the same pipeline in AWS Data Pipeline?

A. You can use multiple compute resources on the same pipeline by defining multiple cluster objects in your definition file and associating the cluster to use for each actMty via its runsOn field.
B. You can use multiple compute resources on the same pipeline by defining multiple cluster definition files.
C. You can use multiple compute resources on the same pipeline by defining multiple clusters for your actMty.
D. You cannot use multiple compute resources on the same pipelin

Answer: A

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

NEW QUESTION 23
True or False: In Amazon EIastiCache replication groups of Redis, for performance tuning reasons, you can change the roles of the cache nodes within the replication group, with the primary and one of the replicas exchanging roles.

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

Answer: C

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

NEW QUESTION 26
How many cg1.4xIarge 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 cg1.4xIarge, 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 28
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 29
You want to define permissions for a role in an IAM policy. Which of the following configuration formats should you use?

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

Answer: C

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

NEW QUESTION 31
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 33
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 36
Which of the following components of AWS Data Pipeline specifies the business logic of your data management?

A. Task Runner
B. Pipeline definition
C. AWS Direct Connect
D. Amazon Simple Storage Service (Amazon S3)

Answer: B

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

NEW QUESTION 38
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 39
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 43
The CFO of a company wants to allow one of his employees to view only the AWS usage report page. Which of the below mentioned IAM policy statements allows the user to have access to the AWS usage report page?

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

Answer: C

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

NEW QUESTION 47
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 52
An organization is planning to setup a management network on the AWS VPC. The organization is trying to secure the webserver on a single VPC instance such that it allows the internet traffic as well as the back-end management traffic. The organization wants to make so that the back end management network

interface can receive the SSH traffic only from a selected IP range, while the internet facing webserver will have an IP address which can receive traffic from all the internet IPs.

How can the organization achieve this by running web server on a single instance?

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

Answer: C

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

NEW QUESTION 55
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 57
An organization is purchasing licensed software. The software license can be registered only to a specific MAC Address. The organization is going to host the software in the AWS environment. How can the organization fulfil the license requirement as the MAC address changes every time an instance is started/stopped/terminated?

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

Answer: C

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

NEW QUESTION 59
Cognito Sync is an AWS service that you can use to synchronize user profile data across mobile devices without requiring your own backend. When the device is online, you can synchronize data. If you also set up push sync, what does it allow you to do?

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

Answer: C

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

NEW QUESTION 64
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 65
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 66
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 70
If a single condition within an IAM policy includes multiple values for one key, it will be evaluated using a logical .

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

Answer: A

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

NEW QUESTION 75
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 76
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 77
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 79
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 82
What RAID method is used on the Cloud Block Storage back-end to implement a very high level of reliability and performance?

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

Answer: C

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

NEW QUESTION 85
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 87
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 90
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 91
An organization has created multiple components of a single application for compartmentalization. Currently all the components are hosted on a single EC2 instance. Due to security reasons the organization wants to implement two separate SSLs for the separate modules although it is already using VPC. How can the organization achieve this with a single instance?

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

Answer: B

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

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

NEW QUESTION 94
How does in-memory caching improve the performance of applications in E|astiCache?

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

Answer: D

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

NEW QUESTION 95
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 100
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 101
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 102
Identify an application that polls AWS Data Pipeline for tasks and then performs those tasks.

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

Answer: C

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

NEW QUESTION 106
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/Iatest/UserGuide/reference_policies_eIements.html

NEW QUESTION 107
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.htmI

NEW QUESTION 112
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 114
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 116
In the context of IAM roles for Amazon EC2, which of the following NOT true about delegating permission to make API requests?

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

Answer: A

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

NEW QUESTION 121
Which of following IAM policy elements lets you specify an exception to a list of actions?

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

Answer: D

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

NEW QUESTION 124
In AWS IAM, which of the following predefined policy condition keys checks how long ago (in seconds) the MFA-validated security credentials making the request were issued using multi-factor authentication
(MFA)?

A. aws:MuItiFactorAuthAge
B. aws:Mu|tiFactorAuthLast
C. aws:MFAAge
D. aws:MuItiFactorAuthPrevious

Answer: A

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

NEW QUESTION 129
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 130
What is the role of the PoIIForTask 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, PoIIForTask 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 134
Who is responsible for modifying the routing tables and networking ACLs in a VPC to ensure that a DB instance is reachable from other instances in the VPC?

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

Answer: B

Explanation:
You are in charge of configuring the routing tables of your VPC as well as the network ACLs rules needed to make your DB instances accessible from all the instances of your VPC that need to communicate with it.
Reference: http://aws.amazon.com/rds/faqs/

NEW QUESTION 139
What is the network performance offered by the c4.8xIarge 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 143
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 147
Do you need to use Amazon Cognito to use the Amazon Mobile Analytics service?

A. N
B. However, it is recommend by AWS to use Amazon Cognito for security best practices.
C. Ye
D. You need to use it only if you have IAM root access.
E. N
F. You cannot use it at all, and you need to use AWS IAM accounts.
G. Ye
H. It is recommended by AWS to use Amazon Cognito to use Amazon Mobile Analytics servic

Answer: A

Explanation:
You can initialize Amazon Mobile Analytics using AWS IAM accounts. AWS recommend using Amazon Cognito for security best practices.
Reference: http://aws.amazon.com/mobi|eanaIytics/faqs/

NEW QUESTION 151
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.htmI

NEW QUESTION 155
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 160
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 165
Which of the following cannot be used to manage Amazon EIastiCache and perform administrative tasks?

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

Answer: D

Explanation:
CIoudWatch is a monitoring tool and doesn't give users access to manage Amazon EIastiCache. Reference:
http://docs.aws.amazon.com/AmazonEIastiCache/latest/UserGuide/WhatIs.NIanaging.htmI

NEW QUESTION 167
Which of the following statements is correct about AWS Direct Connect?

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

Answer: B

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

NEW QUESTION 169
An organization is hosting a scalable web application using AWS. The organization has configured internet facing ELB and Auto Scaling to make the application scalable. Which of the below mentioned
statements is required to be followed when the application is planning to host a web application on VPC?

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

Answer: C

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

NEW QUESTION 174
Is there any way to own a direct connection to Amazon Web Services?

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

Answer: D

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

NEW QUESTION 175
Identify a true statement about the statement ID (Sid) in IAM.

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

Answer: A

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

NEW QUESTION 178
Your company has recently extended its datacenter into a VPC on AVVS to add burst computing capacity as needed Members of your Network Operations Center need to be able to go to the AWS Management Console and administer Amazon EC2 instances as necessary You don't want to create new IAM users for each NOC member and make those users sign in again to the AWS Management Console Which option below will meet the needs for your NOC members?

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

Answer: D

NEW QUESTION 181
You're running an application on-premises due to its dependency on non-x86 hardware and want to use AWS for data backup. Your backup application is only able to write to POSIX-compatible block-based storage. You have 140TB of data and would like to mount it as a single folder on your file server Users must be able to access portions of this data while the backups are taking place. What backup solution would be most appropriate for this use case?

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

Answer: A

NEW QUESTION 185
Your startup wants to implement an order fulfillment process for selling a personalized gadget that needs an average of 3-4 days to produce with some orders taking up to 6 months you expect 10 orders per day on your first day. 1000 orders per day after 6 months and 10,000 orders after 12 months.
Orders coming in are checked for consistency men dispatched to your manufacturing plant for production quality control packaging shipment and payment processing If the product does not meet the quality standards at any stage of the process employees may force the process to repeat a step Customers are notified via email about order status and any critical issues with their orders such as payment failure.
Your case architecture includes AWS Elastic Beanstalk for your website with an RDS MySQL instance for customer data and orders.
How can you implement the order fulfillment process while making sure that the emails are delivered reliably?

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

Answer: C

NEW QUESTION 186
You are tasked with moving a legacy application from a virtual machine running Inside your datacenter to an Amazon VPC Unfortunately this app requires access to a number of on-premises services and no one who configured the app still works for your company. Even worse there's no documentation for it.
What will allow the application running inside the VPC to reach back and access its internal dependencies
without being reconfigured? (Choose 3 answers)

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

Answer: ADF

NEW QUESTION 190
You are the new IT architect in a company that operates a mobile sleep tracking application.
When activated at night, the mobile app is sending collected data points of 1 kilobyte every 5 minutes to your backend.
The backend takes care of authenticating the user and writing the data points into an Amazon DynamoDB table.
Every morning, you scan the table to extract and aggregate last night's data on a per user basis, and store the results in Amazon S3. Users are notified via Amazon SNS mobile push notifications that new data is available, which is parsed and visualized by the mobile app.
Currently you have around 100k users who are mostly based out of North America. You have been tasked to optimize the architecture of the backend system to lower cost. What would you recommend? Choose 2 answers

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

Answer: AD

NEW QUESTION 193
Your department creates regular analytics reports from your company's log files All log data is collected in Amazon S3 and processed by daily Amazon Elastic MapReduce (EMR) jobs that generate daily PDF reports and aggregated tables in CSV format for an Amazon Redshift data warehouse.
Your CFO requests that you optimize the cost structure for this system.
Which of the following alternatives will lower costs without compromising average performance of the system or data integrity for the raw data?

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

Answer: C

NEW QUESTION 197
Your customer wishes to deploy an enterprise application to AWS which will consist of several web servers, several application servers and a small (50GB) Oracle database information is stored, both in the database and the file systems of the various servers. The backup system must support database recovery whole server and whole disk restores, and indMdual file restores with a recovery time of no more than two hours. They have chosen to use RDS Oracle as the database
Which backup architecture will meet these requirements?

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

using traditional enterprise backup software to provide file level restore
D. Backup RDS database to S3 using Oracle RMAN Backup the EC2 instances using Amis, and supplement with EBS snapshots for indMdual volume restore.

Answer: A

NEW QUESTION 200
Your company has HQ in Tokyo and branch offices all over the world and is using a logistics software with a multi-regional deployment on AWS in Japan, Europe and US

A. The logistic software has a 3-tierarchitecture and currently uses MySQL 5.6 for data persistenc
B. Each region has deployed its own database In the HQ region you run an hourly batch process reading data from every region to compute cross-regional reports that are sent by email to all offices this batch process must be completed as fast as possible to quickly optimize logistics how do you build the database architecture in order to meet the requirements'?
C. For each regional deployment, use RDS MySQL with a master in the region and a read replica in theHQ region
D. For each regional deployment, use NIySQL on EC2 with a master in the region and send hourly EBS snapshots to the HQ region
E. For each regional deployment, use RDS MySQL with a master in the region and send hourly RDS snapshots to the HQ region
F. For each regional deployment, use MySQL on EC2 with a master in the region and use S3 to copy data files hourly to the HQ region
G. Use Direct Connect to connect all regional MySQL deployments to the HQ region and reduce network latency for the batch process

Answer: A

NEW QUESTION 205
A web design company currently runs several FTP servers that their 250 customers use to upload and download large graphic files They wish to move this system to AWS to make it more scalable, but they wish to maintain customer privacy and Keep costs to a minimum.
What AWS architecture would you recommend?

A. ASK their customers to use an S3 client instead of an FTP clien
B. Create a single S3 bucket Create an IAM user for each customer Put the IAM Users in a Group that has an IAM policy that permits access to sub-directories within the bucket via use of the 'username' Policy variable.
C. Create a single S3 bucket with Reduced Redundancy Storage turned on and ask their customers to use an S3 client instead of an FTP client Create a bucket for each customer with a Bucket Policy that permits access only to that one customer.
D. Create an auto-scaling group of FTP servers with a scaling policy to automatically scale-in when minimum network traffic on the auto-scaling group is below a given threshol
E. Load a central list of ftp users from S3 as part of the user Data startup script on each Instance.
F. Create a single S3 bucket with Requester Pays turned on and ask their customers to use an S3 client instead of an FTP client Create a bucket tor each customer with a Bucket Policy that permits access only to that one customer.

Answer: A

NEW QUESTION 206
Your company currently has a 2-tier web application running in an on-premises data center. You have experienced several infrastructure failures in the past two months resulting in significant financial losses. Your CIO is strongly agreeing to move the application to AWS. While working on achieving buy-in from the other company executives, he asks you to develop a disaster recovery plan to help improve Business continuity in the short term. He specifies a target Recovery Time Objective (RTO) of 4 hours and a Recovery Point Objective (RPO) of 1 hour or less. He also asks you to implement the solution within 2 weeks. Your database is 200GB in size and you have a 20Mbps Internet connection. How would you do this while minimizing costs?

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

Answer: A

NEW QUESTION 208
You are designing an intrusion detection prevention (IDS/IPS) solution for a customer web application in a single VPC. You are considering the options for implementing IOS IPS protection for traffic coming from the Internet.
Which of the following options would you consider? (Choose 2 answers)

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

Answer: BD

NEW QUESTION 211
You are designing a social media site and are considering how to mitigate distributed denial-of-service (DDoS) attacks. Which of the below are viable mitigation techniques? (Choose 3 answers)

A. Add multiple elastic network interfaces (ENIs) to each EC2 instance to increase the network bandwidth.
B. Use dedicated instances to ensure that each instance has the maximum performance possible.
C. Use an Amazon C|oudFront distribution for both static and dynamic content.

D. Use an Elastic Load Balancer with auto scaling groups at the we
E. App and Amazon Relational Database Service (RDS) tiers
F. Add alert Amazon CloudWatch to look for high Network in and CPU utilization.
G. Create processes and capabilities to quickly add and remove rules to the instance OS firewall

Answer: CEF

NEW QUESTION 216
A web company is looking to implement an external payment service into their highly available application deployed in a VPC Their application EC2 instances are behind a public lacing ELB Auto scaling is used to add additional instances as traffic increases under normal load the application runs 2 instances in the Auto Scaling group but at peak it can scale 3x in size. The application instances need to communicate with the payment service over the Internet which requires whitelisting of all public IP addresses used to communicate with it. A maximum of 4 whitelisting IP addresses are allowed at a time and can be added through an API.
How should they architect their solution?

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

Answer: D

NEW QUESTION 217
Your website is serving on-demand training videos to your workforce. Videos are uploaded monthly in high resolution MP4 format. Your workforce is distributed globally often on the move and using company-provided tablets that require the HTTP Live Streaming (HLS) protocol to watch a video. Your company has no video transcoding expertise and it required you may need to pay for a consultant.
How do you implement the most cost-efficient architecture without compromising high availability and
quality of video delivery'?

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

Answer: C

NEW QUESTION 218
You are running a news website in the eu-west-1 region that updates every 15 minutes. The website has a world-wide audience it uses an Auto Scaling group behind an Elastic Load Balancer and an Amazon
RDS database Static content resides on Amazon S3, and is distributed through Amazon CloudFront. Your Auto Scaling group is set to trigger a scale up event at 60% CPU utilization, you use an Amazon RDSextra large DB instance with 10.000 Provisioned IOPS its CPU utilization is around 80%. While freeable memory is in the 2 GB range.
Web analytics reports show that the average load time of your web pages is around 1.5 to 2 seconds, but your SEO consultant wants to bring down the average load time to under 0.5 seconds.
How would you improve page load times for your users? (Choose 3 answers)

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

Answer: ABD

NEW QUESTION 220
You require the ability to analyze a customer's clickstream data on a website so they can do behavioral analysis. Your customer needs to know what sequence of pages and ads their customer clicked on. This data will be used in real time to modify the page layouts as customers click through the site to increase stickiness and advertising click-through. Which option meets the requirements for captioning and analyzing this data?

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

Answer: B

NEW QUESTION 223
You are designing a connectMty solution between on-premises infrastructure and Amazon VPC. Your servers on-premises will be communicating with your VPC instances. You will be establishing IPSec tunnels over the Internet You will be using VPN gateways, and terminating the IPSec tunnels on AWS supported

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

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

Answer: CDEF

NEW QUESTION 225

Dave is the main administrator in Example Corp., and he decides to use paths to help delineate the users in the company and set up a separate administrator group for each path-based dMsion. Following is a subset of the full list of paths he plans to use:
. /marketing
. /sales
.Hegal
Dave creates an administrator group for the marketing part of the company and calls it NIarketing_Admin. He assigns it the /marketing path. The group's ARN is arn:aws:iam::123456789012:group/marketing/NIarketing_Admin.
Dave assigns the following policy to the NIarketing_Admin group that gives the group permission to use all IAM actions with all groups and users in the /marketing path. The policy also gives the IV|arketing_Admin group permission to perform any AWS S3 actions on the objects in the portion of the corporate bucket.
{
"Version": "2012-10-I7",
"Statement": [
{
"Effect": "Deny",
"Action": "iam:*", "Resource": [
"arn:aws:iam::123456789012:group/marketing/*", "arn:aws:iam::123456789012:user/marketing/*"
I
},
{
"Effect": "A||ow",
"Action": "s3:*",
"Resource": "arn:aws:s3:::exampIe_bucket/marketing/*"
},
{
"Effect": "A||ow", "Action": "s3:ListBucket*",
"Resource": "arn:aws:s3:::exampIe_bucket", "Condition":{"StringLike":{"s3:prefix": "marketing/*"}} I
I I

A. True
B. False

Answer: B

NEW QUESTION 229

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

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

Answer: A

NEW QUESTION 234

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

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

Answer: A

NEW QUESTION 237

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

What database implementation would better fit this scenario, keeping costs as low as possible?

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

Answer: C

NEW QUESTION 241
You are designing a multi-platform web application for AWS The application will run on EC2 instances and will be accessed from PCs. tablets and smart phones Supported accessing platforms are Windows, MacOS, IOS and Android Separate sticky session and SSL certificate setups are required for different platform types which of the following describes the most cost effective and performance efficient architecture setup?

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

Answer: D

NEW QUESTION 246
An administrator is using Amazon CIoudFormation to deploy a three tier web application that consists of a web tier and application tier that will utilize Amazon DynamoDB for storage when creating the CIoudFormation template which of the following would allow the application instance access to the DynamoDB tables without exposing API credentials?

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

Answer: C

NEW QUESTION 250
You have an application running on an EC2 Instance which will allow users to download flies from a private S3 bucket using a pre-signed URL. Before generating the URL the application should verify the existence of the file in S3.
How should the application use AWS credentials to access the S3 bucket securely?

A. Use the AWS account access Keys the application retrieves the credentials from the source code of the application.
B. Create an IAM user for the application with permissions that allow list access to the S3 bucket launch the instance as the IANI user and retrieve the IAM user's credentials from the EC2 instance user data.
C. Create an IAM role for EC2 that allows list access to objects in the S3 bucke
D. Launch the instance with the role, and retrieve the roIe's credentials from the EC2 Instance metadata
E. Create an IAM user for the application with permissions that allow list access to the S3 bucke
F. The application retrieves the IAM user credentials from a temporary directory with permissions that allow read access only to the application user.

Answer: C

NEW QUESTION 254
You are developing a new mobile application and are considering storing user preferences in AWS.2w This would provide a more uniform cross-device experience to users using multiple mobile devices to access the application. The preference data for each user is estimated to be 50KB in size Additionally 5 million customers are expected to use the application on a regular basis. The solution needs to be
cost-effective, highly available, scalable and secure, how would you design a solution to meet the above requirements?

A. Setup an RDS MySQL instance in 2 availability zones to store the user preference dat
B. Deploy apublic facing application on a server in front of the database to manage security and access credentials
C. Setup a DynamoDB table with an item for each user having the necessary attributes to hold the user preference
D. The mobile application will query the user preferences directly from the DynamoDB tabl
E. Utilize ST
F. Web Identity Federation, and DynamoDB Fine Grained Access Control to authenticate and authorize access.
G. Setup an RDS MySQL instance with multiple read replicas in 2 availability zones to store the user preference data .The mobile application will query the user preferences from the read replica
H. Leverage the MySQL user management and access prMlege system to manage security and access credentials.
I. Store the user preference data in S3 Setup a DynamoDB table with an item for each user and an item attribute pointing to the user' S3 objec
J. The mobile application will retrieve the S3 URL from DynamoDB and then access the S3 object directly utilize STS, Web identity Federation, and S3 ACLs to authenticate and authorize access.

Answer: B

NEW QUESTION 256
A web-startup runs its very successful social news application on Amazon EC2 with an Elastic Load Balancer, an Auto-Scaling group of Java/Tomcat application-servers, and DynamoDB as data store. The main web-application best runs on m2 x large instances since it is highly memory- bound Each new deployment requires semi-automated creation and testing of a new AMI for the application servers which takes quite a while ana is therefore only done once per week.
Recently, a new chat feature has been implemented in nodejs and wails to be integrated in the architecture. First tests show that the new component is CPU bound Because the company has some experience with using Chef, they decided to streamline the deployment process and use AWS Ops Works as an application life cycle tool to simplify management of the application and reduce the deployment cycles.
What configuration in AWS Ops Works is necessary to integrate the new chat module in the most cost-efficient and filexible way?

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

Answer: C

NEW QUESTION 261
In AWS, which security aspects are the customer's responsibility? Choose 4 answers

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

Answer: ACDF

NEW QUESTION 265
Your company policies require encryption of sensitive data at rest. You are considering the possible options for protecting data while storing it at rest on an EBS data volume, attached to an EC2 instance. Which of these options would allow you to encrypt your data at rest? Choose 3 answers

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

Answer: ACD

NEW QUESTION 269
You have recently joined a startup company building sensors to measure street noise and air quality in urban areas. The company has been running a pilot deployment of around 100 sensors for 3 months each sensor uploads 1KB of sensor data every minute to a backend hosted on AWS.
During the pilot, you measured a peak or 10 IOPS on the database, and you stored an =average of 3GB of sensor data per month in the database.
The current deployment consists of a load-balanced auto scaled Ingestion layer using EC2 instances and a PostgreSQL RDS database with 500GB standard storage.
The pilot is considered a success and your CEO has managed to get the attention or some potential investors. The business plan requires a deployment of at least 100K sensors which needs to be supported by the backend. You also need to store sensor data for at least two years to be able to compare year over year Improvements.
To secure funding, you have to make sure that the platform meets these requirements and leaves room for further scaling. Which setup win meet the requirements?

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

Answer: C

NEW QUESTION 271
A company is storing data on Amazon Simple Storage Service (S3). The company's security policy mandates that data is encrypted at rest. Which of the following methods can achieve this?
Choose 3 answers

A. Use Amazon S3 server-side encryption with AWS Key Management Service managed keys.
B. Use Amazon S3 server-side encryption with customer-provided keys.
C. Use Amazon S3 server-side encryption with EC2 key pair.
D. Use Amazon S3 bucket policies to restrict access to the data at rest.
E. Encrypt the data on the client-side before ingesting to Amazon S3 using their own master key.
F. Use SSL to encrypt the data while in transit to Amazon S3.

Answer: ABE

NEW QUESTION 275
You've been hired to enhance the overall security posture for a very large e-commerce site They have a well architected multi-tier application running in a VPC that uses ELBs in front of both the web and the app tier with static assets served directly from S3 They are using a combination of RDS and DynamoOB for their dynamic data and then archMng nightly into S3 for further processing with EMR They are concerned because they found questionable log entries and suspect

someone is attempting to gain unauthorized access.
Which approach provides a cost effective scalable mitigation to this kind of attack?

A. Recommend that they lease space at a DirectConnect partner location and establish a 1G DirectConnect connection to their VPC they would then establish Internet connectMty into their space, filter the traffic in hardware Web Application Firewall (WAF). And then pass the traffic through the DirectConnect connection into their application running in their VPC.
B. Add previously identified hostile source IPs as an explicit INBOUND DENY NACL to the web tier subnet
C. Add a WAF tier by creating a new ELB and an AutoScaling group of EC2 Instances running ahost-based WAF They would redirect Route 53 to resolve to the new WAF tier ELB The WAF tier would their pass the traffic to the current web tier The web tier Security Groups would be updated to only allow traffic from the WAF tier Security Group
D. Remove all but TLS 1.2 from the web tier ELB and enable Advanced Protocol Filtering This will enable the ELB itself to perform WAF functionality.

Answer: C

NEW QUESTION 278
Your company has an on-premises multi-tier PHP web application, which recently experienced downtime due to a large burst In web traffic due to a company announcement Over the coming days, you are expecting similar announcements to drive similar unpredictable bursts, and are looking to find ways to quickly improve your infrastructures ability to handle unexpected increases in traffic.
The application currently consists of 2 tiers a web tier which consists of a load balancer and several Linux Apache web servers as well as a database tier which hosts a Linux server hosting a MySQL database. Which scenario below will provide full site functionality, while helping to improve the ability of your application in the short timeframe required?

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

Answer: C

NEW QUESTION 283
You are designing the network infrastructure for an application server in Amazon VPC. Users will access all application instances from the Internet, as well as from an on-premises network. The on-premises
network is connected to your VPC over an AWS Direct Connect link. How would you design routing to meet the above requirements?

A. Configure a single routing table with a default route via the Internet gatewa
B. Propagate a default route via BGP on the AWS Direct Connect customer route
C. Associate the routing table with all VPC subnets.
D. Configure a single routing table with a default route via the Internet gatewa
E. Propagate specific routes for the on-premises networks via BGP on the AWS Direct Connect customer route
F. Associate the routing table with all VPC subnets.
G. Configure a single routing table with two default routes: on to the Internet via an Internet gateway, the other to the on-premises network via the VPN gatewa
H. Use this routing table across all subnets in the VPC.
I. Configure two routing tables: on that has a default router via the Internet gateway, and other that has a default route via the VPN gatewa
J. Associate both routing tables with each VPC subnet.

Answer: A

NEW QUESTION 286
Auto Scaling requests are signed with a signature calculated from the request and the user's private key.

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

Answer: C

NEW QUESTION 290
What does elasticity mean to AWS?

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

Answer: B

NEW QUESTION 294
How is AWS readily distinguished from other vendors in the traditional IT computing landscape?

A. Experience
B. Scalable and elasti

C. Secur
D. Cost-effectiv
E. Reliable
F. Secur
G. Flexibl
H. Cost-effectiv
I. Scalable and elasti
J. Global
K. Secur
L. Flexibl
M. Cost-effectiv
N. Scalable and elasti
O. Experienced
P. Flexibl
Q. Cost-effectiv
R. Dynami
S. Secur
T. Experience

Answer: C

NEW QUESTION 297
Your company is storing millions of sensitive transactions across thousands of 100-GB files that must be encrypted in transit and at rest. Analysts concurrently depend on subsets of files, which can consume up to 5 TB of space, to generate simulations that can be used to steer business decisions. You are required to design an AWS solution that can cost effectively accommodate the long-term storage and in-flight subsets of data.

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

Answer: D

NEW QUESTION 298
Your customer is willing to consolidate their log streams (access logs application logs security logs etc.) in one single system. Once consolidated, the customer wants to analyze these logs in real time based on heuristics. From time to time, the customer needs to validate heuristics, which requires going back to data samples extracted from the last 12 hours?
What is the best approach to meet your customer's requirements?

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

Answer: B

NEW QUESTION 301
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

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

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

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

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

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

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