

Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional



NEW QUESTION 1

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

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

Answer: D

Explanation:

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

Reference: <http://aws.amazon.com/cognito/faqs/>

NEW QUESTION 2

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

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

Answer: B

Explanation:

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

Reference:

<http://docs.aws.amazon.com/AmazonS3/latest/dev/access-control-auth-workflow-object-operation.html>

NEW QUESTION 3

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

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

Answer: B

Explanation:

The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. The organization wants to extend their network into the cloud and also directly access the internet from their AWS VPC. Thus, the organization should setup a Virtual Private Cloud (VPC) with a public subnet and a private subnet, and a virtual private gateway to enable communication with their data center network over an IPsec VPN tunnel. To setup this configuration the organization needs to use the Amazon VPC with a VPN connection. The organization network administrator must designate a physical appliance as a customer gateway and configure it. The organization would need the below mentioned information to setup this configuration:

The type of customer gateway, such as Cisco ASA, Juniper J-Series, Juniper SSG, Yamaha
Internet-routable IP address (static) of the customer gateway's external interface

Border Gateway Protocol (BGP) Autonomous System Number (ASN) of the customer gateway, if the organization is creating a dynamically routed VPN connection.

Internal network IP ranges that the user wants to advertise over the VPN connection to the VPC. Reference:

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

NEW QUESTION 4

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

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

Answer: B

Explanation:

The user passwords (login profiles) of IAM users can contain any Basic Latin (ASCII) characters. Reference:

<http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html>

NEW QUESTION 5

An organization is planning to host a Wordpress blog as well a Joomla CMS on a single instance launched with VPC. The organization wants to have separate domains for each application and assign them using Route 53. The organization may have about ten instances each with two applications as mentioned above.

While launching the instance, the organization configured two separate network interfaces (primary + ENI) and wanted to have two elastic IPs for that instance. It was suggested to use a public IP from AWS instead of an elastic IP as the number of elastic IPs is restricted. What action will you recommend to the organization?

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

Answer: B

Explanation:

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

NEW QUESTION 6

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

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

Answer: A

Explanation:

The maximum number of VPCs allowed per region is 5.

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

NEW QUESTION 7

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

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

Answer: D

Explanation:

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

Thus, when the user wants to reject traffic from the selected IPs it is recommended to use ACL with subnets.

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

NEW QUESTION 8

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

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

Answer: D

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. The VPC also work with IAM and the organization can create IAM users who have access to various VPC services. The organization can setup access for the IAM user who can modify the security groups of the VPC. The sample policy is given below:

```
{
  "Version": "2012-10-17",
  "Statement":
  [
    {
      "Effect": "Allow",
      "Action": "ec2:RunInstances",
      "Resource":
      [
        "arn:aws:ec2:region::image/ami-*",
        "arn:aws:ec2:region:account:subnet/subnet-1a2b3c4d",
        "arn:aws:ec2:region:account:network-interface/*",
        "arn:aws:ec2:region:account:volume/*",
        "arn:aws:ec2:region:account:key-pair/*",
        "arn:aws:ec2:region:account:security-group/sg-123abc123"
      ]
    }
  ]
}
```

With this policy the user can create four subnets in separate zones and provide IAM user access to each subnet

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

NEW QUESTION 9

An organization is planning to host an application on the AWS VPC. The organization wants dedicated instances. However, an AWS consultant advised the

organization not to use dedicated instances with VPC as the design has a few limitations. Which of the below mentioned statements is not a limitation of dedicated instances with VPC?

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

Answer: C

Explanation:

The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Dedicated instances are Amazon EC2 instances that run in a Virtual Private Cloud (VPC) on hardware that is dedicated to a single customer. The client's dedicated instances are physically isolated at the host hardware level from instances that are not dedicated instances as well as from instances that belong to other AWS accounts.

All instances launched with the dedicated tenancy model of VPC will always be dedicated instances. Dedicated tenancy has a limitation that it may not support a few services, such as RDS. Even the EBS will not be on dedicated hardware. However the user can save some cost as well as reserve some capacity by using a Reserved Instance model with dedicated tenancy.

Reference: <http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/dedicated-instance.html>

NEW QUESTION 10

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

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

Answer: C

Explanation:

In Amazon IAM, the maximum length for a role name is 64 characters.

Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html>

NEW QUESTION 10

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 12

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

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

Answer: C

Explanation:

Generally AWS EC2 allows running 20 on-demand instances and 100 spot instances at a time. This limit can be increased by requesting at <https://aws.amazon.com/contact-us/ec2-request>. Excluding certain types of instances, the limit is lower than mentioned above. For g2.2xlarge, the user can run only 5

on-demand instance at a time.

Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_ec2

NEW QUESTION 14

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

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

Answer: C

Explanation:

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

Reference: <http://docs.aws.amazon.com/cognito/devguide/identity/concepts/authentication-flow/>

NEW QUESTION 18

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 subnet
- 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 group
- D. Launch RDS with that subnet group.
- E. Create a network interface and attach two subnets to it
- 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 19

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 21

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

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

Answer: A

Explanation:

In Amazon VPC, you can assign any Private IP address to your instance as long as it is: Part of the associated subnet's IP address range

Not reserved by Amazon for IP networking purposes Not currently assigned to another interface Reference: <http://aws.amazon.com/vpc/faqs/>

NEW QUESTION 26

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

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

Answer: C

Explanation:

If an explicit deny is not found among the applicable policies for a specific request, IAM's Policy Evaluation Logic checks for any "allow" instructions to check if the request can be successfully completed.

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

NEW QUESTION 27

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

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

Answer: B

Explanation:

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

NEW QUESTION 29

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:

Multiple compute resources can be used on the same pipeline in AWS Data Pipeline by defining multiple cluster objects in your definition file and associating the cluster to use for each 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/datapipeline/faqs/>

NEW QUESTION 32

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

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

Answer: D

Explanation:

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

NEW QUESTION 33

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

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

Answer: B

Explanation:

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

NEW QUESTION 36

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

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

Answer: D

Explanation:

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

Reference: <http://aws.amazon.com/ec2/instance-types/>

NEW QUESTION 41

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/SNSMobilePush.html>

NEW QUESTION 43

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

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

Answer: B

Explanation:

AWS has many solutions for Disaster Recovery(DR) and High Availability(HA). When the organization wants to have HA and DR but are okay to have a longer recovery time they should select the option backup and restore with S3. The data can be sent to S3 using either Direct Connect, Storage Gateway or over the internet.

The EC2 instance will pick the data from the S3 bucket when started and setup the environment. This process takes longer but is very cost effective due to the low pricing of S3. In all the other options, the EC2 instance might be running or there will be AMI storage costs.

Thus, it will be a costlier option. In this scenario the organization should plan appropriate tools to take a backup, plan the retention policy for data and setup security of the data.

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

NEW QUESTION 47

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

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

Answer: A

Explanation:

In Amazon ElastiCache, you can run a maximum of 20 Cache Nodes. Reference: <http://aws.amazon.com/elasticache/faqs/>

NEW QUESTION 50

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

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

Answer: C

Explanation:

An AWS Direct Connect location provides access to Amazon Web Services in the region it is associated with, as well as access to other US regions. For example, you can provision a single connection to any AWS Direct Connect location in the US and use it to access public AWS services in all US Regions and AWS GovCloud (US).

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

NEW QUESTION 51

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

- A. Task Runner
- B. Pipeline definition

- C. AWS Direct Connect
- D. Amazon Simple Storage Service (Amazon S3)

Answer: B

Explanation:

A pipeline definition specifies the business logic of your data management.

Reference: <http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/what-is-datapipeline.html>

NEW QUESTION 54

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 58

What types of identities do Amazon Cognito identity pools support?

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

Answer: A

Explanation:

Amazon Cognito identity pools support both authenticated and unauthenticated identities. Authenticated identities belong to users who are authenticated by a public login provider or your own backend authentication process. Unauthenticated identities typically belong to guest users. Reference:

<http://docs.aws.amazon.com/cognito/devguide/identity/identity-pools/>

NEW QUESTION 59

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

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

Answer: C

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the CFO wants to allow only AWS usage report page access, the policy for that IAM user will be as given below:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow", "Action": [
        "aws-portal:ViewUsage"
      ],
      "Resource": "*"
    }
  ]
}
```

Reference: <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/billing-permissions-ref.html>

NEW QUESTION 62

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

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

Answer: B

Explanation:

The maximum number of BGP advertised routes allowed per route table is 100.

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

NEW QUESTION 66

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 69

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

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

Answer: B

Explanation:

IOPS are input/output operations per second. Amazon EBS measures each I/O operation per second (that is 256 KB or smaller) as one IOPS.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-io-characteristics.html>

NEW QUESTION 72

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

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

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

Answer: C

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. An Elastic Network Interface (ENI) is a virtual network interface that the user can attach to an instance in a VPC.

The user can create a management network using two separate network interfaces. For the present scenario it is required that the secondary network interface on the instance handles the public facing traffic and the primary network interface handles the back-end management traffic and it is connected to a separate subnet in the VPC that has more restrictive access controls. The public facing interface, which may or may not be behind a load balancer, has an associated security group to allow access to the server from the internet while the private facing interface has an associated security group allowing SSH access only from an allowed range of IP addresses either within the VPC or from the internet, a private subnet within the VPC or a virtual private gateway.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html>

NEW QUESTION 75

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 80

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.html>

NEW QUESTION 83

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

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

Answer: B

Explanation:

The maximum length for an instance profile name is 128 characters.

Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html>

NEW QUESTION 84

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

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

Answer: C

Explanation:

A Provisioned IOPS (SSD) volume can range in size from 4 GiB to 16 TiB and you can provision up to 20,000 IOPS per volume.

Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html#EBSVolumeTypes_piops

NEW QUESTION 88

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 90

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 93

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

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

Answer: A

Explanation:

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

NEW QUESTION 96

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

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

Answer: B

Explanation:

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

NEW QUESTION 99

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 100

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 ephemeral 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 ephemeral storage for cost savings.

Reference: <http://aws.amazon.com/ec2/purchasing-options/spot-instances/>

NEW QUESTION 105

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

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

Answer: A

Explanation:

To send push notifications to mobile devices using Amazon SNS and ADM, you need to obtain the following: Registration ID and Client secret.

Reference: <http://docs.aws.amazon.com/sns/latest/dg/SNSMobilePushPrereq.html>

NEW QUESTION 109

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

- . All the EC2 instances should have a private IP
- . All the EC2 instances should receive data via the ELB's. Which of these will not be needed in this setup?

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

Answer: A

Explanation:

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

- . route all inbound traffic from ELB to EC2 instances
- . route all outbound traffic from EC2 instances through NAT

Reference: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/AWSHowTo-vpc.html>

NEW QUESTION 113

An EC2 instance that performs source/destination checks by default is launched in a private VPC subnet. All security, NACL, and routing definitions are configured as expected. A custom NAT instance is launched.

Which of the following must be done for the custom NAT instance to work?

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

Answer: A

Explanation:

Each EC2 instance performs source/destination checks by default. This means that the instance must be the source or destination of any traffic it sends or receives. However, a NAT instance must be able to send and receive traffic when the source or destination is not itself. Therefore, you must disable source/destination checks on the NAT instance.

Reference:

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_NAT_Instance.html#EIP_Disable_SrcDestCheck

NEW QUESTION 117

An organization is making software for the CIA in US

- A. CIA agreed to host the application on AWS but in a secure environment
- B. The organization is thinking of hosting the application on the AWS GovCloud region
- C. Which of the below mentioned differences is not correct when the organization is hosting on the AWS GovCloud in comparison with the AWS standard region?
- D. The billing for the AWS GovCloud will be in a different account than the Standard AWS account.
- E. GovCloud region authentication is isolated from Amazon.com.
- F. Physical and logical administrative access only to U.S. persons.
- G. persons.

H. It is physically isolated and has logical network isolation from all the other region

Answer: A

Explanation:

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

NEW QUESTION 118

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

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

Answer: D

Explanation:

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

NEW QUESTION 120

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

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

Answer: B

Explanation:

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

NEW QUESTION 122

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

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

Answer: C

Explanation:

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

NEW QUESTION 127

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

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

Answer: B

Explanation:

You can provision up to 3TB storage and 30,000 IOPS per database instance. For a workload with 50% writes and 50% reads running on a `cr1.8xlarge` instance, you can realize over 25,000 IOPS for PostgreSQL. However, by provisioning more than this limit, you may be able to achieve lower latency and higher throughput.

Your actual realized IOPS may vary from the amount you provisioned based on your database workload, instance type, and database engine choice.
Reference: <https://aws.amazon.com/rds/postgresql/>

NEW QUESTION 130

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

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

Answer: C

Explanation:

AWS Data Pipeline is a web service that helps you reliably process and move data between different AWS compute and storage services as well as on-premise data sources at specified intervals. With AWS Data Pipeline, you can regularly access your data where it's stored, transform and process it at scale, and efficiently transfer the results to another AWS.

AWS Data Pipeline helps you easily create complex data processing workloads that are fault tolerant, repeatable, and highly available. AWS Data Pipeline also allows you to move and process data that was previously locked up in on-premise data silos. Reference: <http://aws.amazon.com/datapipeline/>

NEW QUESTION 131

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

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

Answer: C

Explanation:

A task runner is an application that polls AWS Data Pipeline for tasks and then performs those tasks. You can either use Task Runner as provided by AWS Data Pipeline, or create a custom Task Runner application.

Task Runner is a default implementation of a task runner that is provided by AWS Data Pipeline. When Task Runner is installed and configured, it polls AWS Data Pipeline for tasks associated with pipelines that you have activated. When a task is assigned to Task Runner, it performs that task and reports its status back to AWS Data Pipeline. If your workflow requires non-default behavior, you'll need to implement that functionality in a custom task runner.

Reference:

<http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/dp-how-remote-taskrunner-client.html>

NEW QUESTION 136

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

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

Answer: C

Explanation:

Within an IAM policy, IfExists can be added to the end of any condition operator except the Null condition. It can be used to indicate that conditional comparison needs to happen if the policy key is present in the context of a request; otherwise, it can be ignored.

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

NEW QUESTION 140

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

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

Answer: A

Explanation:

A service account is a special Google account that can be used by applications to access Google services programmatically. This account belongs to your application or a virtual machine (VM), instead of to an individual end user. Your application uses the service account to call the Google API of a service, so that the users aren't directly involved.

A service account can have zero or more pairs of service account keys, which are used to authenticate to Google. A service account key is a public/private keypair generated by Google. Google retains the public key, while the user is given the private key.

Reference: <https://cloud.google.com/iam/docs/service-accounts>

NEW QUESTION 143

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

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

Answer: C

Explanation:

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

NEW QUESTION 146

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

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

Answer: A

Explanation:

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

NEW QUESTION 150

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

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

Answer: A

Explanation:

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

NEW QUESTION 152

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

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

Answer: C

Explanation:

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

NEW QUESTION 157

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

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

Answer: D

Explanation:

AWS Direct Connect is a network service that provides an alternative to using the Internet to utilize AWS cloud service. Hence, a Direct Connect link cannot be connected to the Internet directly.

Reference: <http://aws.amazon.com/directconnect/faqs/>

NEW QUESTION 159

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

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

Answer: D

Explanation:

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

NEW QUESTION 161

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

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

Answer: A

Explanation:

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

NEW QUESTION 162

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

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

Answer: A

Explanation:

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

NEW QUESTION 166

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

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

Answer: C

Explanation:

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

NEW QUESTION 168

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

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

Answer: A

Explanation:

Your application does not need to reside in the same VPC as the CloudHSM instance.

However, the server or instance on which your application and the HSM client is running must have network (IP) reachability to the HSM. You can establish network connectivity in a variety of ways, including operating your application in the same VPC, with VPC peering, with a VPN connection, or with Direct Connect.

Reference: <https://aws.amazon.com/cloudhsm/faqs/>

NEW QUESTION 170

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

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

Answer: A

Explanation:

Amazon ElastiCache cache security groups are only applicable to cache clusters that are not running in an Amazon Virtual Private Cloud environment (VPC). If you are running in an Amazon Virtual Private Cloud, Cache Security Groups is not available in the console navigation pane.

Reference: <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/CacheSecurityGroup.html>

NEW QUESTION 171

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

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

Answer: C

Explanation:

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

NEW QUESTION 174

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

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

Answer: B

Explanation:

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

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-workload-demand.html>

NEW QUESTION 179

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 181

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

- A. The organization should plan the app server on the public subnet and database in the organization's data center and connect them with the VPN gateway.

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

Answer: A

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account.

The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to

connect with his data centre. When the user has configured this setup with Wizard, it will create a virtual private gateway to route all the traffic of the VPN subnet.

If the virtual private gateway is attached with VPC and the user deletes the VPC from the console it will first automatically detach the gateway and only then delete the VPC.

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

NEW QUESTION 184

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

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

Answer: C

Explanation:

A Provisioned IOPS (SSD) volume can range in size from 4 GiB to 16 TiB and you can provision up to 20,000 IOPS per volume. The ratio of IOPS provisioned to the volume size requested should be a maximum of 30; for example, a volume with 3000 IOPS must be atleast 100 GB.

Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html#EBSVolumeTypes_piops

NEW QUESTION 189

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

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

Answer: D

Explanation:

Provisioned IOPS volumes are designed to meet the needs of I/O-intensive workloads, particularly database workloads that are sensitive to storage performance and consistency in random access I/O throughput. A provisioned IOPS volume can range in size from 10 GB to 1 TB and the user can provision up to 4000 IOPS per volume.

The ratio of IOPS provisioned to the volume size requested can be a maximum of 30; for example, a volume with 3000 IOPS must be atleast 100 GB.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html>

NEW QUESTION 194

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

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

Answer: A

Explanation:

The two reasons you would need to edit claims issued in a SAML token are: The NameIdentifier claim cannot be the same as the username stored in AD, and The app requires a different set of claim URIs.

Reference:

<https://azure.microsoft.com/en-us/documentation/articles/active-directory-saml-claims-customization/>

NEW QUESTION 199

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

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

Answer: D

Explanation:

AWS CloudHSM provides secure cryptographic key storage to customers by making hardware security modules (HSMs) available in the AWS cloud. AWS CloudHSM requires the following environment before an HSM appliance can be provisioned. A virtual private cloud (VPC) in the region where you want the AWS CloudHSM service.

- One private subnet (a subnet with no Internet gateway) in the VPC. The HSM appliance is provisioned into this subnet.
- One public subnet (a subnet with an Internet gateway attached). The control instances are attached to this subnet.
- An AWS Identity and Access Management (IAM) role that delegates access to your AWS resources to AWS CloudHSM.
- An EC2 instance, in the same VPC as the HSM appliance, that has the SafeNet client software installed. This instance is referred to as the control instance and is used to connect to and manage the HSM appliance.
- A security group that has port 22 (for SSH) or port 3389 (for RDP) open to your network. This security group is attached to your control instances so you can access them remotely.

NEW QUESTION 203

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

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

Answer: D

Explanation:

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

NEW QUESTION 204

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

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

Answer: A

Explanation:

You can initialize Amazon Mobile Analytics using AWS IAM accounts. AWS recommends using Amazon Cognito for security best practices.
 Reference: <http://aws.amazon.com/mobileanalytics/faqs/>

NEW QUESTION 209

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

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

Answer: D

Explanation:

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

NEW QUESTION 210

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

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

Answer: C

Explanation:

AWS security follows the shared security model where the user is as much responsible as Amazon. Since Amazon is a public cloud it is bound to be targeted by hackers. If an organization is planning to host their application on AWS EC2, they should perform the below mentioned security checks as a measure to find any

security weakness/data leaks:

Perform penetration testing as performed by attackers to find any vulnerability. The organization must take an approval from AWS before performing penetration testing

Perform hardening testing to find if there are any unnecessary ports open Perform SQL injection to find any DB security issues

The code memory checks are generally useful when the organization wants to improve the application performance.

Reference: <http://aws.amazon.com/security/penetration-testing/>

NEW QUESTION 213

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

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

Answer: D

Explanation:

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

Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

NEW QUESTION 218

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

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

Answer: B

Explanation:

The VPC must allow non-dedicated EC2 instances. You cannot use ElastiCache in a VPC that is configured for dedicated instance tenancy.

Reference: <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/AmazonVPC.EC.html>

NEW QUESTION 220

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

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

Answer: C

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources, such as RDS into a virtual network that the user has defined. Subnets are segments of a VPC's IP address range that the user can designate to a group of VPC resources based on security and operational needs. A DB subnet group is a collection of subnets (generally private) that the user can create in a VPC and which the user assigns to the RDS DB instances. A DB subnet group allows the user to specify a particular VPC when creating DB instances. If the RDS instance is required to be accessible from the internet:

The organization must setup that the RDS instance is enabled with the VPC attributes, DNS hostnames and DNS resolution.

The organization must enable the parameter in the console which makes the RDS instance publicly accessible.

The organization must allow access from the internet in the RDS VPC security group. Reference:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.html

NEW QUESTION 224

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

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

Answer: D

NEW QUESTION 228

To serve Web traffic for a popular product your chief financial officer and IT director have purchased 10 ml large heavy utilization Reserved Instances (RIs) evenly

spread across two availability zones: Route 53 is used to deliver the traffic to an Elastic Load Balancer (ELB). After several months, the product grows even more popular and you need additional capacity. As a result, your company purchases two C3.2xlarge medium utilization instances. You register the two c3.2xlarge instances with your ELB and quickly find that the m1.large instances are at 100% of capacity and the c3.2xlarge instances have significant capacity that's unused. Which option is the most cost effective and uses EC2 capacity most effectively?

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

Answer: B

NEW QUESTION 233

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

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

Answer: BE

NEW QUESTION 237

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 242

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

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

Answer: D

NEW QUESTION 244

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 249

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

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

Answer: A

NEW QUESTION 253

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

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

Answer: C

NEW QUESTION 258

An AWS customer is deploying an application that is composed of an AutoScaling group of EC2 Instances.

The customer's security policy requires that every outbound connection from these instances to any other service within the customer's Virtual Private Cloud must be authenticated using a unique X.509 certificate that contains the specific instance-id.

In addition, X.509 certificates must be designed by the customer's Key management service in order to be trusted for authentication.

Which of the following configurations will support these requirements?

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

Answer: A

NEW QUESTION 260

Your company runs a customer-facing event registration site. This site is built with a 3-tier architecture with web and application tier servers and a MySQL database. The application requires 6 web tier servers and 6 application tier servers for normal operation, but can run on a minimum of 65% server capacity and a single MySQL database. When deploying this application in a region with three availability zones (AZs) which architecture provides high availability?

- A. A web tier deployed across 2 AZs with 3 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer), and an application tier deployed across 2 AZs with 3 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and one RDS (Relational Database Service) instance deployed with read replicas in the other AZ.
- B. A web tier deployed across 3 AZs with 2 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer) and an application tier deployed across 3 AZs with 2 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and one RDS (Relational Database Service) instance deployed with read replicas in the two other AZs.
- C. A web tier deployed across 2 AZs with 3 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer) and an application tier deployed across 2 AZs with 3 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and a Multi-AZ RDS (Relational Database Service) deployment.
- D. A web tier deployed across 3 AZs with 2 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer). And an application tier deployed across 3 AZs with 2 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and a Multi-AZ RDS (Relational Database Service) deployment.

Answer: D

NEW QUESTION 265

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 application.
- B. Develop a CloudFormation 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 267

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

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

Answer: C

NEW QUESTION 268

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

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

Answer: A

NEW QUESTION 269

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

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

Answer: AB

NEW QUESTION 271

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

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

Answer: AC

NEW QUESTION 274

You must architect the migration of a web application to AWS. The application consists of Linux web servers running a custom web server. You are required to save the logs generated from the application to a durable location. What options could you select to migrate the application to AWS? (Choose 2)

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

Answer: AD

NEW QUESTION 277

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 queue
- B. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few days
- 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 days
- 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 days
- 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 queue
- K. S3 to host videos with Lifecycle Management to archive all files to Glacier after a few days
- L. CloudFront to serve HLS transcoded videos from Glacier.

Answer: C

NEW QUESTION 279

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

Which architecture change could you introduce to reduce costs and still keep your web application secure and scalable?

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

Answer: C

NEW QUESTION 281

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 capturing 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 286

You have deployed a three-tier web application in a VPC with a CIDR block of 10.0.0.0/28 You initially deploy two web servers, two application servers, two

database servers and one NAT instance for a total of seven EC2 instances. The web application and database servers are deployed across two availability zones (AZs). You also deploy an ELB in front of the two web servers, and use Route53 for DNS. Web traffic gradually increases in the first few days following the deployment, so you attempt to double the number of instances in each tier of the application to handle the new load. Unfortunately, some of these new instances fail to launch.

Which of the following could be the root cause? (Choose 2 answers)

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

Answer: CE

NEW QUESTION 289

You are designing an SaaS solution that requires HTTPS clients to be authenticated by the Webserver using client certificate authentication. The solution must be resilient.

Which of the following options would you consider for configuring the web server infrastructure? (Choose 2 answers)

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

Answer: AB

NEW QUESTION 292

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

You have a regulatory requirement that all data needs to be encrypted before being uploaded to the cloud. How do you implement this in a highly available and cost-efficient way?

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

Answer: C

NEW QUESTION 293

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

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

- A. End-to-end protection of data in transit
- B. End-to-end Identity authentication
- C. Data encryption 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 294

Your Fortune 500 company has undertaken a TCO analysis evaluating the use of Amazon S3 versus acquiring more hardware. The outcome was that all employees would be granted access to use Amazon S3 for storage of their personal documents.

Which of the following will you need to consider so you can set up a solution that incorporates single sign-on from your corporate AD or LDAP directory and restricts access for each user to a designated user folder in a bucket? (Choose 3 Answers)

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

Answer: ABD

NEW QUESTION 299

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 it 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 302

Your company hosts a social media site supporting users in multiple countries. You have been asked to provide a highly available design for the application that leverages multiple regions for the most recently accessed content and latency sensitive portions of the web site. The most latency sensitive component of the application involves reading user preferences to support web site personalization and ad selection. In addition to running your application in multiple regions, which option will support this application's requirements?

- A. Serve user content from S3. CloudFront and use Route53 latency-based routing between ELBs in each region. Retrieve user preferences from a local DynamoDB table in each region and leverage SQS to capture changes to user preferences with SOS workers for propagating updates to each table.
- B. Use the S3 Copy API to copy recently accessed content to multiple regions and serve user content from S3. CloudFront with dynamic content and an ELB in each region. Retrieve user preferences from an ElasticCache cluster in each region and leverage SNS notifications to propagate user preference changes to a worker node in each region.
- C. Use the S3 Copy API to copy recently accessed content to multiple regions and serve user content from S3. CloudFront and Route53 latency-based routing between ELBs. In each region, retrieve user preferences from a DynamoDB table and leverage SQS to capture changes to user preferences with SOS workers for propagating DynamoDB updates.
- D. Serve user content from S3. CloudFront with dynamic content, and an ELB in each region. Retrieve user preferences from an ElasticCache cluster in each region and leverage Simple Workflow (SWF) to manage the propagation of user preferences from a centralized OB to each ElasticCache cluster.

Answer: A

NEW QUESTION 303

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

- A. Use OAuth 2.0 to retrieve temporary AWS security credentials to enable your NOC members to sign in to the AWS Management Console.
- B. Use web Identity Federation to retrieve AWS temporary security credentials to enable your NOC members to sign in to the AWS Management Console.
- C. Use your on-premises SAML 2.0-compliant identity provider (IDP) to grant the NOC members federated access to the AWS Management Console via the AWS single sign-on (SSO) endpoint.
- D. Use your on-premises SAML2.0-compliant identity provider (IDP) to retrieve temporary security credentials to enable NOC members to sign in to the AWS Management Console.

Answer: D

NEW QUESTION 307

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions 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](#)