

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

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



NEW QUESTION 1

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 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 3

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

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

Answer: A

Explanation:

The maximum number of VPCs allowed per region is 5.

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

NEW QUESTION 4

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

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

Answer: D

Explanation:

AWS provides an on demand, scalable infrastructure. AWS EC2 allows the user to launch On-Demand instances and the organization should create an AMI of the running instance. When the organization is experiencing varying loads and the time of the load is not known but it is higher than the routine traffic it is recommended that the organization launches a few instances before hand and then setups AutoScaling with policies which scale up and down as per the EC2 metrics, such as Network I/O or CPU utilization.

If the organization keeps all 10 additional instances as a part of the AutoScaling policy sometimes during a sudden higher load it may take time to launch instances and may not give an optimal performance. This is the reason it is recommended that the organization keeps an additional 5 instances running and the next 5 instances scheduled as per the AutoScaling policy for cost effectiveness.

Reference: http://media.amazonwebservices.com/AWS_Web_Hosting_Best_Practices.pdf

NEW QUESTION 5

An organization 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 6

In which step of using AWS Direct Connect should the user determine the required port speed?

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

Answer: D

Explanation:

To submit an AWS Direct Connect connection request, you need to provide the following information: Your contact information. The AWS Direct Connect Location to connect to.

Details of AWS Direct Connect partner if you use the AWS Partner Network (APN) service. The port speed you require, either 1 Gbps or 10 Gbps.

Reference: <http://docs.aws.amazon.com/directconnect/latest/UserGuide/getstarted.html#ConnectionRequest>

NEW QUESTION 7

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

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

Answer: C

Explanation:

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

NEW QUESTION 8

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 9

A user has created a MySQL RDS instance with PIOPS. Which of the below mentioned statements will help user understand the advantage of PIOPS?

- A. The user can achieve additional dedicated capacity for the EBS I/O with an enhanced RDS option
- B. It uses a standard EBS volume with optimized configuration the stacks
- C. It uses optimized EBS volumes and optimized configuration stacks
- D. It provides a dedicated network bandwidth between EBS and RDS

Answer:

C

Explanation:

RDS DB instance storage comes in two types: standard and provisioned IOPS. Standard storage is allocated on the Amazon EBS volumes and connected to the user's DB instance. Provisioned IOPS uses optimized EBS volumes and an optimized configuration stack. It provides additional, dedicated capacity for the EBS I/O.
Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html>

NEW QUESTION 10

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

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

Answer: A**Explanation:**

The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances.

There are two ELBs available with VPC: internet facing and internal (private) ELB. For internal servers, such as App servers the organization can create an internal load balancer in their VPC and then place back-end application instances behind the internal load balancer. The internal load balancer will route requests to the back-end application instances, which are also using private IP addresses and only accept requests from the internal load balancer.

The Internal ELB supports only one subnet in each AZ and asks the user to select a subnet while configuring internal ELB.

Reference: http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/USVPC_creating_basic_lb.html

NEW QUESTION 10

IV|apMySite is setting up a web application in the AWS VPC. The organization has decided to use an AWS RDS instead of using its own DB instance for HA and DR requirements.

The organization also wants to secure RDS access. How should the web application be setup with RDS?

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

Answer: B**Explanation:**

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources, such as RDS into a virtual network that the user has defined. Subnets are segments of a VPC's IP address range that the user can designate to a group of VPC resources based on the security and operational needs.

A DB subnet group is a collection of subnets (generally private) that a user can create in a VPC and assign to the RDS DB instances. A DB subnet group allows the user to specify a particular VPC when creating the DB instances. Each DB subnet group should have subnets in at least two Availability Zones in a given region.

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

NEW QUESTION 11

What bandwidths do AWS Direct Connect currently support?

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

Answer: D**Explanation:**

AWS Direct Connection currently supports 1Gbps and 10 Gbps.

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

NEW QUESTION 12

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

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

Answer: A

Explanation:

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

NEW QUESTION 15

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 18

An organization (account ID 123412341234) has configured the IAM policy to allow the user to modify his credentials. What will the below mentioned statement allow the user to perform?

```
{  
  "Version": "2012-10-17",  
  "Statement": [{  
    "Effect": "Allow", "Action": [ "iam:AddUserToGroup",  
    "iam:RemoveUserFromGroup", "iam:GetGroup"  
  ]  
  },  
  {  
    "Resource": "arn:aws:iam:: 123412341234:group/TestingGroup"  
  }  
}]
```

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

Answer: A

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the organization (account ID 123412341234) wants their users to manage their subscription to the groups, they should create a relevant policy for that. The below mentioned policy allows the respective IAM user to update the membership of the group called MarketingGroup.

```
{  
  "Version": "2012-10-17",  
  "Statement": [{  
    "Effect": "Allow", "Action": [ "iam:AddUserToGroup",  
    "iam:RemoveUserFromGroup", "iam:GetGroup"  
  ]  
  },  
  {  
    "Resource": "arn:aws:iam:: 123412341234:group/ TestingGroup "  
  }  
}]
```

Reference:

<http://docs.aws.amazon.com/IAM/latest/UserGuide/Credentials-Permissions-examples.htm#creds-policy-credentials>

NEW QUESTION 21

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

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

Answer: D

Explanation:

In AWS Data Pipeline, an activity fails if all of its activity attempts return with a failed state. By default, an activity retries three times before entering a hard failure state. You can increase the number of automatic retries to 10. However, the system does not allow indefinite retries.

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

NEW QUESTION 22

How many c4.xlarge on-demand instances can a user run in one region without taking any limit increase approval from AWS?

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

Answer: B

Explanation:

Generally AWS EC2 allows running 20 on-demand instances and 100 spot instances at a time. This limit can be increased by requesting at <https://aws.amazon.com/contact-us/ec2-request>. Excluding certain types of instances, the limit is lower than mentioned above. For cg1.4xlarge, the user can run only 2 on-demand instances at a time.
Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_ec2

NEW QUESTION 24

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

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

Answer: D

Explanation:

The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances.

There are two ELBs available with VPC: internet facing and internal (private) ELB. For internal servers, such as App servers the organization can create an internal load balancer in their VPC and then place back-end application instances behind the internal load balancer. The internal load balancer will route requests to the back-end application instances, which are also using private IP addresses and only accept requests from the internal load balancer.

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/vpc-loadbalancer-types.html>

NEW QUESTION 28

True or False: Amazon ElastiCache supports the Redis key-value store.

- A. True, ElastiCache supports the Redis key-value store, but with limited functionalities.
- B. False, ElastiCache does not support the Redis key-value store.
- C. True, ElastiCache supports the Redis key-value store.
- D. False, ElastiCache supports the Redis key-value store only if you are in a VPC environment.

Answer: C

Explanation:

This is true. ElastiCache supports two open-source in-memory caching engines: 1. Memcached - a widely adopted memory object caching system. ElastiCache is protocol compliant with Memcached, so popular tools that you use today with existing Memcached environments will work seamlessly with the service. 2. Redis - a popular open-source in-memory key-value store that supports data structures such as sorted sets and lists. ElastiCache supports Master / Slave replication and Multi-AZ which can be used to achieve cross AZ redundancy.

Reference: <https://aws.amazon.com/elasticache/>

NEW QUESTION 30

Which of the following is NOT an advantage of using AWS Direct Connect?

- A. AWS Direct Connect provides users access to public and private resources by using two different connections while maintaining network separation between the public and private environments.
- B. AWS Direct Connect provides a more consistent network experience than Internet-based connections.
- C. AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS.
- D. AWS Direct Connect reduces your network cost.

Answer: A

Explanation:

AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your datacenter, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections.

By using industry standard 802.1q VLANs, this dedicated connection can be partitioned into multiple virtual interfaces. This allows you to use the same connection to access public resources such as objects stored in Amazon S3 using public IP address space, and private resources such as Amazon EC2 instances running within an Amazon Virtual Private Cloud (VPC) using private IP space, while maintaining network separation between the public and private environments.

Reference: <http://aws.amazon.com/directconnect/#details>

NEW QUESTION 33

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

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

Answer: B

Explanation:

You can use the action Modify DB Instance, available in the Amazon RDS API, to pass values for the parameters DB Instance Identifier and DB Security Groups

specifying the instance ID and the DB Security Groups you want your instance to be part of.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/APIReference/API_ModifyDBInstance.html

NEW QUESTION 36

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 38

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 43

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 48

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 53

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

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

Answer: A

Explanation:

Temporary credentials in IAM are valid throughout their defined duration of time and hence can't be revoked. However, because permissions are evaluated each time an AWS request is made using the credentials, you can achieve the effect of revoking the credentials by changing the permissions for the credentials even after they have been issued. Reference:

http://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp_control-access_disable-perms.html

NEW QUESTION 56

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 59

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 60

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 63

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 66

An organization is undergoing a security audit. The auditor wants to view the AWS VPC configurations as the organization has hosted all the applications in the AWS VPC. The auditor is from a remote place and wants to have access to AWS to view all the VPC records.

How can the organization meet the expectations of the auditor without compromising on the security of their AWS infrastructure?

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

Answer: C

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. The VPC also works with IAM and the organization can create IAM users who have access to various VPC services.

If an auditor wants to have access to the AWS VPC to verify the rules, the organization should be careful before sharing any data which can allow making updates to the AWS infrastructure. In this scenario it is recommended that the organization creates an IAM user who will have read only access to the VPC. Share the above mentioned credentials with the auditor as it cannot harm the organization. The sample policy is given below:

```
{
  "Effect": "Allow",
  "Action": [ "ec2:DescribeVpcs", "ec2:DescribeSubnets",
    "ec2:DescribeInternetGateways", "ec2:DescribeCustomerGateways", "ec2:DescribeVpnGateways", "ec2:DescribeVpnConnections", "ec2:DescribeRouteTables",
    "ec2:DescribeAddresses", "ec2:DescribeSecurityGroups", "ec2:DescribeNetworkAcls", "ec2:DescribeDhcpOptions", "ec2:DescribeTags", "ec2:DescribeInstances"
  ],
  "Resource": "*"
}
```

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

NEW QUESTION 71

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 74

An organization is planning to create a secure scalable application with AWS VPC and ELB. The organization has two instances already running and each instance has an ENI attached to it in addition to a primary network interface. The primary network interface and additional ENI both have an elastic IP attached to it.

If those instances are registered with ELB and the organization wants ELB to send data to a particular EIP of the instance, how can they achieve this?

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

Answer: A

Explanation:

Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances. There are two ELBs available with VPC: internet facing and internal (private) ELB. For the internet facing ELB it is required that the ELB should be in a public subnet.

When the user registers a multi-homed instance (an instance that has an Elastic Network Interface (ENI) attached) with a load balancer, the load balancer will route the traffic to the IP address of the primary network interface (eth0).

Reference: <http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/gs-ec2VPC.html>

NEW QUESTION 78

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

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

Answer: D

Explanation:

The maximum length for a certificate ID is 128 characters.

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

NEW QUESTION 82

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 83

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

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

Answer: C

Explanation:

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

NEW QUESTION 86

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

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

Answer: B

Explanation:

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

NEW QUESTION 90

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 92

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

- A. Allow only IAM users to connect with the EC2 instances with their own secret access key.
- B. Create a procedure to revoke the access rights of the individual user when they are not required to connect to EC2 instance anymore for the purpose of

application configuration.

- C. Apply the latest patch of OS and always keep it updated.
- D. Disable the password based login for all the user
- E. All the users should use their own keys to connect with the instance securely.

Answer: A

Explanation:

Since AWS is a public cloud any application hosted on EC2 is prone to hacker attacks. It becomes extremely important for a user to setup a proper security mechanism on the EC2 instances. A few of the security measures are listed below:

Always keep the OS updated with the latest patch

Always create separate users with in OS if they need to connect with the EC2 instances, create their keys and disable their password

Create a procedure using which the admin can revoke the access of the user when the business work on the EC2 instance is completed

Lock down unnecessary ports

Audit any proprietary applications that the user may be running on the EC2 instance

Provide temporary escalated prMleges, such as sudo for users who need to perform occasional prMleged tasks

The IAM is useful when users are required to work with AWS resources and actions, such as launching an instance. It is not useful to connect (RDP / SSH) with an instance.

Reference: <http://aws.amazon.com/articles/1233/>

NEW QUESTION 95

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 99

True or False : "In the context of Amazon ElastiCache, from the application's point of view, connecting to the cluster configuration endpoint is no different than connecting directly to an indMdual cache node."

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

Answer: B

Explanation:

This is true. From the application's point of view, connecting to the cluster configuration endpoint is no different than connecting directly to an indMdual cache node. In the process of connecting to cache nodes, the application resolves the configuration endpoint's DNS name. Because the configuration endpoint maintains CNAME entries for all of the cache nodes, the DNS name resolves to one of the nodes; the client can then connect to that node.

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

NEW QUESTION 103

A user is thinking to use EBS PIOPS volume. Which of the below mentioned options is a right use case for the PIOPS EBS volume?

- A. Analytics
- B. System boot volume
- C. Nlongo DB
- D. Log processing

Answer: C

Explanation:

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

NEW QUESTION 105

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 106

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 111

When using string conditions within IAM, short versions of the available comparators can be used instead of the more verbose ones. `streql` is the short version of the string condition.

- A. `StringEqualsIgnoreCase`
- B. `StringNotEqualsIgnoreCase`
- C. `StringLikeStringEquals`
- D. `StringNotEquals`

Answer: A

Explanation:

When using string conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions. For instance, `streql` is the short version of `StringEqualsIgnoreCase` that checks for the exact match between two strings ignoring their case.

Reference: <http://awsdocs.s3.amazonaws.com/SNS/20100331/sns-gsg-2010-03-31.pdf>

NEW QUESTION 114

Select the correct statement about Amazon ElastiCache.

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

Answer: A

Explanation:

ElastiCache is a web service that makes it easy to set up, manage, and scale a distributed in-memory cache environment in the cloud. It provides a high-performance, scalable, and cost-effective caching solution, while removing the complexity associated with deploying and managing a distributed cache environment. With ElastiCache, you can quickly deploy your cache environment, without having to provision hardware or install software.

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

NEW QUESTION 117

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 120

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 123

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

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

Answer: C

Explanation:

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

NEW QUESTION 126

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

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

Answer: B

Explanation:

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

NEW QUESTION 129

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 131

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 136

An organization has developed an application which provides a smarter shopping experience. They need to show a demonstration to various stakeholders who may not be able to access the in premise application so they decide to host a demo version of the application on AWS. Consequently they will need a fixed elastic IP attached automatically to the instance when it is launched.

In this scenario which of the below mentioned options will not help assign the elastic IP automatically?

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 141

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

- A. ExamKiller has to setup one centralized VPC which will peer in to all the other VPCs of the tenants.
- B. ExamKiller should setup VPC peering with all the VPCs peering each other but block the IPs from CIDR of the tenant VPCs to deny them.
- C. ExamKiller should setup all the VPCs with the same CIDR but have a centralized VP
- D. This way only the centralized VPC can talk to the other VPCs using VPC peering.
- E. ExamKiller should setup all the VPCs meshed together with VPC peering for all VPC

Answer: A

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. A VPC peering connection allows the user to route traffic between the peer VPCs using private IP addresses as if they are a part of the same network.

This is helpful when one VPC from the same or different AWS account wants to connect with resources of the other VPC. The organization wants to setup that one VPC can connect with all the other VPCs but all other VPCs cannot connect among each other. This can be achieved by configuring VPC peering where one VPC is peered with all the other VPCs, but the other VPCs are not peered to each other. The VPCs are in the same or a separate AWS account and should not have overlapping CIDR blocks.

Reference:

<http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/peering-configurations-full-access.html#many-vpcs-full-access>

NEW QUESTION 145

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

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

Answer: C

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. The user can create a subnet with the same size of VPC. However, he cannot create any other subnet since the CIDR of the second subnet will conflict with the first subnet.

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

NEW QUESTION 149

Identify a true statement about using an IAM role to grant permissions to applications running on Amazon EC2 instances.

- A. When AWS credentials are rotated, developers have to update only the root Amazon EC2 instance that uses their credentials.
- B. When AWS credentials are rotated, developers have to update only the Amazon EC2 instance on which the password policy was applied and which uses their credentials.
- C. When AWS credentials are rotated, you don't have to manage credentials and you don't have to worry about long-term security risks.
- D. When AWS credentials are rotated, you must manage credentials and you should consider precautions for long-term security risks.

Answer: C

Explanation:

Using IAM roles to grant permissions to applications that run on EC2 instances requires a bit of extra configuration. Because role credentials are temporary and rotated automatically, you don't have to manage credentials, and you don't have to worry about long-term security risks.

Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/role-usecase-ec2app.html>

NEW QUESTION 150

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 153

In the context of Amazon ElastiCache CLI, which of the following commands can you use to view all ElastiCache instance events for the past 24 hours?

- A. `elasticache-events --duration 24`
- B. `elasticache-events --duration 1440`
- C. `elasticache-describe-events --duration 24`
- D. `elasticache describe-events --source-type cache-cluster --duration 1440`

Answer: D

Explanation:

In Amazon ElastiCache, the code `"aws elasticache describe-events --source-type cache-cluster --duration 1440"` is used to list the cache-cluster events for the past 24 hours (1440 minutes). Reference: <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/ECEvents.Viewing.html>

NEW QUESTION 157

When using Numeric Conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions. Which of the following is the short version of the Numeric Condition "NumericLessThanEquals"?

- A. `numlteq`
- B. `numlteql`
- C. `numltequals`
- D. `numeql`

Answer: A

Explanation:

When using Numeric Conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions. For instance, `numlteq` is the short version of `NumericLessThanEquals`.

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

NEW QUESTION 160

AWS has launched T2 instances which come with CPU usage credit. An organization has a requirement which keeps an instance running for 24 hours. However, the organization has high usage only during 11 AM to 12 PM. The organization is planning to use a T2 small instance for this purpose. If the organization already has multiple instances running since Jan 2012, which of the below mentioned options should the organization implement while launching a T2 instance?

- A. The organization must migrate to the EC2-VPC platform first before launching a T2 instance.
- B. While launching a T2 instance the organization must create a new AWS account as this account does not have the EC2-VPC platform.
- C. Create a VPC and launch a T2 instance as part of one of the subnets of that VPC.
- D. While launching a T2 instance the organization must select EC2-VPC as the platform.

Answer: C

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. The AWS account provides two platforms:

EC2-CLASSIC and EC2-VPC, depending on when the user has created his AWS account and which regions he is using. If the user has created the AWS account

after 2013-12-04, it supports only EC2-VPC. In this scenario, since the account is before the required date the supported platform will be EC2-CLASSIC. It is required that the organization creates a VPC as the T2 instances can be launched only as a part of VPC.
Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/vpc-migrate.html>

NEW QUESTION 162

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

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

Answer: D

Explanation:

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

NEW QUESTION 165

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 170

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

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

Answer: A

Explanation:

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

NEW QUESTION 175

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 177

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 181

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 183

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 185

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 188

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

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

Answer: A

Explanation:

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

NEW QUESTION 192

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 at least 100 GB.

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

NEW QUESTION 194

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 PIOPS 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 198

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 200

Which of the following components of AWS Data Pipeline polls for tasks and then performs those tasks?

- A. Pipeline Definition
- B. Task Runner
- C. Amazon Elastic MapReduce (EMR)
- D. AWS Direct Connect

Answer: B

Explanation:

Task Runner polls for tasks and then performs those tasks.

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

NEW QUESTION 203

In Amazon ElastiCache, the default cache port is:

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

Answer: D

Explanation:

In Amazon ElastiCache, you can specify a new port number for your cache cluster, which by default is 11211 for Memcached and 6379 for Redis.

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

NEW QUESTION 207

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 212

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

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

Answer: C

Explanation:

The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances. There are two ELBs available with VPC: internet facing and internal (private) ELB. For internet facing ELB it is required that ELB should be in a public subnet.

After the user creates the public subnet, he should ensure to associate the route table of the public subnet with the internet gateway to enable the load balancer in the subnet to connect with the internet. Reference: <http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/CreateVPCForELB.html>

NEW QUESTION 216

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 221

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

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

Answer: A

Explanation:

The Sid(statement ID) is an optional identifier that you provide for the policy statement. You can assign a Sid a value to each statement in a statement array. In IAM, the Sid is not exposed in the IAM API. You can't retrieve a particular statement based on this ID.

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

NEW QUESTION 222

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 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 226

A read only news reporting site with a combined web and application tier and a database tier that receives large and unpredictable traffic demands must be able to respond to these traffic fluctuations automatically. What AWS services should be used meet these requirements?

- A. Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaling group monitored with CloudWatch and RDS with read replicas.
- B. Stateful instances for the web and application tier in an autoscaling group monitored with CloudWatch and RDS with read replicas.
- C. Stateful instances for the web and application tier in an autoscaling group monitored with CloudWatc
- D. And multi-AZ RDS.
- E. Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaling group monitored with CloudWatch and multi-AZ RDS.

Answer: A

NEW QUESTION 229

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 231

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

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

Answer: D

NEW QUESTION 235

You have been asked to design the storage layer for an application. The application requires disk performance of at least 100,000 IOPS. In addition, the storage layer must be able to survive the loss of an indMdual disk, EC2 instance, or Availability Zone without any data loss. The volume you provide must have a capacity of at least 3 TB. Which of the following designs will meet these objectives?

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

Answer: C

NEW QUESTION 238

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 241

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 244

Your department creates regular analytics reports from your company's log files. All log data is collected in Amazon S3 and processed by daily Amazon Elastic MapReduce (EMR) jobs that generate daily PDF reports and aggregated tables in CSV format for an Amazon Redshift data warehouse.

Your CFO requests that you optimize the cost structure for this system.

Which of the following alternatives will lower costs without compromising average performance of the system or data integrity for the raw data?

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

Answer: C

NEW QUESTION 246

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 250

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

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

Answer: A

NEW QUESTION 252

A web design company currently runs several FTP servers that their 250 customers use to upload and download large graphic files. They wish to move this system to AWS to make it more scalable, but they wish to maintain customer privacy and keep costs to a minimum.

What AWS architecture would you recommend?

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

Answer: A

NEW QUESTION 253

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

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

Answer: AC

NEW QUESTION 258

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

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

Answer: A

NEW QUESTION 259

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 VNI 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 the 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 261

An AWS customer runs a public blogging website. The site users upload two million blog entries a month. The average blog entry size is 200 KB. The access rate to blog entries drops to negligible 6 months after publication and users rarely access a blog entry 1 year after publication. Additionally, blog entries have a high update rate during the first 3 months following publication, this drops to no updates after 6 months. The customer wants to use CloudFront to improve his user's load times. Which of the following recommendations would you make to the customer?

- A. Duplicate entries into two different buckets and create two separate CloudFront distributions where S3 access is restricted only to Cloud Front identity
- B. Create a CloudFront distribution with "US Europe" price class for US/Europe users and a different CloudFront distribution with "All Edge Locations" for the remaining users.
- C. Create a CloudFront distribution with S3 access restricted only to the CloudFront identity and partition the blog entry's location in S3 according to the month it was uploaded to be used with CloudFront behaviors.
- D. Create a CloudFront distribution with Restrict Viewer Access Forward Query string set to true and minimum TTL of 0.

Answer: C

NEW QUESTION 265

An International company has deployed a multi-tier web application that relies on DynamoDB in a single region For regulatory reasons they need disaster recovery capability In a separate region with a Recovery Time Objective of 2 hours and a Recovery Point Objective of 24 hours. They should synchronize their data on a regular basis and be able to provision the web application rapidly using CloudFormation.

The objective is to minimize changes to the existing web application, control the throughput of DynamoDB used for the synchronization of data and synchronize only the modified elements.

Which design would you choose to meet these requirements?

- A. Use AWS data Pipeline to schedule a DynamoDB cross region copy once a day, create a "Lastupdated" attribute in your DynamoDB table that would represent the timestamp of the last update and use it as a filter.
- B. Use EMR and write a custom script to retrieve data from DynamoDB in the current region using a SCAN operation and push it to DynamoDB in the second region.
- C. Use AWS data Pipeline to schedule an export of the DynamoDB table to S3 in the current region once a day then schedule another task immediately after it that will import data from S3 to DynamoDB in the other region.
- D. Send also each Ante into an SQS queue in me second region; use an auto-scaling group behind the SQS queue to replay the write in the second region.

Answer: A

NEW QUESTION 269

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

Answer: C

NEW QUESTION 270

A corporate web application is deployed within an Amazon Virtual Private Cloud (VPC) and is connected to the corporate data center via an IPsec VPN. The application must authenticate against the

on-premises LDAP server. After authentication, each logged-in user can only access an Amazon Simple Storage Space (S3) keyspace specific to that user. Which two approaches can satisfy these objectives? (Choose 2 answers)

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

Answer: BC

NEW QUESTION 271

Your company previously configured a heavily used, dynamically routed VPN connection between your on-premises data center and AWS. You recently provisioned a DirectConnect connection and would like to start using the new connection. After configuring DirectConnect settings in the AWS Console, which of the following options win provide the most seamless transition for your users?

- A. Delete your existing VPN connection to avoid routing loops configure your DirectConnect router with the appropriate settings and verify network traffic is leveraging DirectConnect.
- B. Configure your DirectConnect router with a higher BGP priority man your VPN router, verify network traffic is leveraging Directconnect and then delete your existing VPN connection.
- C. Update your VPC route tables to point to the DirectConnect connection configure your DirectConnect router with the appropriate settings verify network traffic is leveraging DirectConnect and then delete the VPN connection.
- D. Configure your DirectConnect router, update your VPC route tables to point to the DirectConnect connection, configure your VPN connection with a higher BGP priorit
- E. And verify network traffic is leveraging the DirectConnect connection.

Answer: D

NEW QUESTION 275

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

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

- A. Replace the Auto Scaling launch configuration to include c3.8xlarge instances; those instances can potentially yield a network throuthput of 10gbps.
- B. Re-architect your ingest pattern, have the app authenticate against your identity provider, and use your identity provider as a broker fetching temporary AWS

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

Answer: C

NEW QUESTION 280

You have deployed a three-tier web application in a VPC with a CIDR block of 10.0.0.0/28 You initially deploy two web servers, two application sewers, two database sewers and one NAT instance tor a total of seven EC2 instances The web. Application and database sewers are deployed across two availability zones (AZs). You also deploy an ELB in front of the two web servers, and use Route53 for DNS Web (raffile gradually increases in the first few days following the deployment, so you attempt to double the number of instances in each tier of the application to handle the new load unfortunately some of these new instances fail to launch.

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

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

Answer: CE

NEW QUESTION 285

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

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

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

- A. Manage encryption keys on-premises in an encrypted relational databas
- B. Set up an on-premises server with sufficient storage to temporarily store files, and then upload them to Amazon S3, providing a client-side master key.
- C. Mange encryption keys in a Hardware Security Module (HSM) appliance on-premises serve r with sufficient storage to temporarily store, encrypt, and upload files directly into Amazon Glacier.
- D. Nlamage 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 CloudHSMN appliance
- F. Encrypt files prior to uploading on the employee desktop, and then upload directly into Amazon Glacier.

Answer: C

NEW QUESTION 288

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

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

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

Answer: CDEF

NEW QUESTION 292

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

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

- A. Add an Elastic Load Balancing health check to the Auto Scaling grou
- B. Set a short period for the health checks to operate as soon as possible in order to prevent premature registration of the instance to theload balancer.
- C. Enable EC2 instance CloudWatch alerts to change the launch configuration's AMI to the previous on
- D. Gradually terminate instances that are using the new AMI.
- E. Set the Elastic Load Balancing health check configuration to target a part of the application that fully tests application health and returns an error if the tests fail.
- F. Create a new launch configuration that refers to the new AMI, and associate it with the grou
- G. Double the size of the group, wait for the new instances to become healthy, and reduce back to the original size.If new instances do not become healthy, associate the previous launch configuration.

H. Increase the Elastic Load Balancing Unhealthy Threshold to a higher value to prevent an unhealthy instance from going into service behind the load balancer.

Answer: CD

NEW QUESTION 294

You are running a successful multitier web application on AWS and your marketing department has asked you to add a reporting tier to the application. The reporting tier will aggregate and publish status reports every 30 minutes from user-generated information that is being stored in your web application's database. You are currently running a Multi-AZ RDS MySQL instance for the database tier. You also have implemented ElastiCache as a database caching layer between the application tier and database tier. Please select the answer that will allow you to successfully implement the reporting tier with as little impact as possible to your database.

- A. Continually send transaction logs from your master database to an S3 bucket and generate the reports off the S3 bucket using S3 byte range requests.
- B. Generate the reports by querying the synchronously replicated standby RDS MySQL instance maintained through Multi-AZ.
- C. Launch a RDS Read Replica connected to your Multi AZ master database and generate reports by querying the Read Replica.
- D. Generate the reports by querying the ElastiCache database caching tier

Answer: C

NEW QUESTION 298

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

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

Answer: B

NEW QUESTION 303

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

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

Answer: C

NEW QUESTION 306

Your company plans to host a large donation website on Amazon Web Services (AWS). You anticipate a large and undetermined amount of traffic that will create many database writes. To be certain that you do not drop any writes to a database hosted on AWS. Which service should you use?

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

Answer: B

NEW QUESTION 307

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 310

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

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

Answer: D

NEW QUESTION 314

You deployed your company website using Elastic Beanstalk and you enabled log file rotation to S3. An Elastic Map Reduce job is periodically analyzing the logs on S3 to build a usage dashboard that you share with your CIO.

You recently improved overall performance of the website using Cloud Front for dynamic content delivery and your website as the origin.

After this architectural change, the usage dashboard shows that the traffic on your website dropped by an order of magnitude. How do you fix your usage dashboard'?

- A. Enable Cloud Front to deliver access logs to S3 and use them as input of the Elastic Map Reduce job.
- B. Turn on Cloud Trail and use trail log tiles on S3 as input of the Elastic Map Reduce job
- C. Change your log collection process to use Cloud Watch ELB metrics as input of the Elastic MapReduce job
- D. Use Elastic Beanstalk "Rebuild Environment" option to update log delivery to the Elastic IV|ap Reduce job.
- E. Use Elastic Beanstalk "Restart App server(s)" option to update log delivery to the Elastic Map Reduce job.

Answer: D

NEW QUESTION 318

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

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

- * AWS-Certified-Solutions-Architect-Professional Questions and Answers Updated Frequently
- * AWS-Certified-Solutions-Architect-Professional Practice Questions Verified by Expert Senior Certified Staff
- * AWS-Certified-Solutions-Architect-Professional Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * AWS-Certified-Solutions-Architect-Professional Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year