

Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional



NEW QUESTION 1

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

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

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 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

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

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

Answer: D

Explanation:

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

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

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

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

NEW QUESTION 4

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 5

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 6

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

- A. ArnCheck
- B. ArnMatch
- C. ArnCase
- D. ArnLike

Answer: D

Explanation:

Amazon Resource Name (ARN) condition operators let you construct Condition elements that restrict access based on comparing a key to an ARN. ArnLike, for instance, is a case-insensitive matching of the ARN. Each of the six colon-delimited components of the ARN is checked separately and each can include a multi-character match wildcard (*) or a single-character match wildcard (?).

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

NEW QUESTION 7

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 8

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.html#creds-policies-credentials>

NEW QUESTION 9

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

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

Answer: B

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. If the organization has multiple VPN connections, he can provide secure communication between sites using the AWS VPN CloudHub.

The VPN CloudHub operates on a simple hub-and-spoke model that the user can use with or without a VPC. This design is suitable for customers with multiple branch offices and existing internet connections who would like to implement a convenient, potentially low-cost hub-and-spoke model for primary or backup connectivity between remote offices.

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

NEW QUESTION 10

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

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

Answer: C

Explanation:

In Amazon ElastiCache, a replication group is a collection of Redis Cache Clusters, with one primary read-write cluster and up to five secondary, read-only clusters, which are called read replicas. You can change the roles of the cache clusters within the replication group, with the primary cluster and one of the replicas exchanging roles. You might decide to do this for performance tuning reasons.

Reference: <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/Replication.Redis.Groups.html>

NEW QUESTION 10

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

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

Answer: A

Explanation:

In Amazon SNS, you have the ability to send notification messages directly to apps on mobile devices. Notification messages sent to a mobile endpoint can appear in the mobile app as message alerts, badge updates, or even sound alerts. Microsoft Windows Mobile Messaging (MWMM) doesn't exist and is not supported by Amazon SNS.

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

NEW QUESTION 11

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 13

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 15

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 18

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 20

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 23

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 27

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

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

Answer: C

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. An Elastic Network Interface (ENI) is a virtual network interface that the user can attach to an instance in a VPC. An ENI can include attributes such as: a primary private IP address, one or more secondary private IP addresses, one elastic IP address per private IP address, one public IP address, one or more security groups, a MAC address, a source/destination check flag, and a description.

The user can create a network interface, attach it to an instance, detach it from an instance, and attach it to another instance. The attributes of a network interface follow the network interface as it is attached or detached from an instance and reattached to another instance. Thus, the user can maintain a fixed MAC using the network interface.

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

NEW QUESTION 32

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 36

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

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

Answer: C

Explanation:

Cognito Sync is an AWS service that you can use to synchronize user profile data across mobile devices without requiring your own backend. When the device is online, you can synchronize data, and if you have

also set up push sync, notify other devices immediately that an update is available. Reference: <http://docs.aws.amazon.com/cognito/devguide/sync/>

NEW QUESTION 37

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 EI

Answer: A

Explanation:

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

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

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

NEW QUESTION 42

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_pio ps

NEW QUESTION 47

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 48

What RAID method is used on the Cloud Block Storage back-end to implement a very high level of reliability and performance?

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

Answer: C

Explanation:

Cloud Block Storage back-end storage volumes employs the RAID 10 method to provide a very high level of reliability and performance.

Reference: http://www.rackspace.com/knowledge_center/product-faq/cloud-block-storage

NEW QUESTION 53

One of the AWS account owners faced a major challenge in June as his account was hacked and the hacker deleted all the data from his AWS account. This resulted in a major blow to the business.

Which of the below mentioned steps would not have helped in preventing this action?

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

Answer: C

Explanation:

AWS security follows the shared security model where the user is as much responsible as Amazon. If the user wants to have secure access to AWS while hosting applications on EC2, the first security rule to follow is to enable MFA for all users. This will add an added security layer. In the second step, the user should never

give his access or secret access keys to anyone as well as store inside programs. The better solution is to use IAM roles. For critical data of the organization, the user should keep an offsite/ in premise backup which will help to recover critical data in case of security breach.

It is recommended to have AWS AMIs and snapshots as well as keep them at other regions so that they will help in the DR scenario. However, in case of a data security breach of the account they may not be very helpful as hacker can delete that.

Therefore, creating an AMI and a snapshot of the data at regular intervals as well as keep a copy to separate regions, would not have helped in preventing this action.

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

NEW QUESTION 54

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

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

Answer: D

Explanation:

In Amazon ElastiCache, in-memory caching improves application performance by storing critical pieces of data in memory for low-latency access. Cached information may include the results of I/O-intensive database queries or the results of computationally intensive calculations.

Reference: <http://aws.amazon.com/elasticache/faqs/#g4>

NEW QUESTION 58

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 60

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

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

Answer: B

Explanation:

Amazon EC2 uses an instance profile as a container for an IAM role. When you create an IAM role using the console, the console creates an instance profile automatically and gives it the same name as the role it corresponds to. If you use the AWS CLI, API, or an AWS SDK to create a role, you create the role and instance profile as separate actions, and you might give them different names.

Reference:

http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2_instance-profiles.html

NEW QUESTION 61

Attempts, one of the three types of items associated with the schedule pipeline in the AWS Data Pipeline, provides robust data management.

Which of the following statements is NOT true about Attempts?

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

Answer: C

Explanation:

Attempts, one of the three types of items associated with a schedule pipeline in AWS Data Pipeline, provides robust data management. AWS Data Pipeline retries a failed operation. It continues to do so until the task reaches the maximum number of allowed retry attempts. Attempt objects track the various attempts, results, and failure reasons if applicable. Essentially, it is the instance with a counter. AWS Data Pipeline performs retries using the same resources from the previous attempts, such as Amazon EMR clusters and EC2 instances.

Reference:

<http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/dp-how-tasks-scheduled.html>

NEW QUESTION 62

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 64

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

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

Answer: C

Explanation:

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

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

Reference:

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

NEW QUESTION 69

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 70

What happens when Dedicated instances are launched into a VPC?

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

Answer: C

Explanation:

If you launch an instance into a VPC that has an instance tenancy of dedicated, your instance is automatically a Dedicated instance, regardless of the tenancy of the instance.

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

NEW QUESTION 73

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

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

E. The organization must create a subnet group with VPC using more than one subnet which are a part of separate AZs.

Answer: A

Explanation:

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

Each DB subnet group should have subnets in at least two Availability Zones in a given region. If the RDS instance is required to be accessible from the internet the organization must enable the VPC attributes, DNS hostnames and DNS resolution. For each RDS DB instance that the user runs in a VPC, he should reserve at least one address in each subnet in the DB subnet group for use by Amazon RDS for recovery actions.

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

NEW QUESTION 77

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 82

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 83

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 87

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 92

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. numeq

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 93

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 98

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 102

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

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

Answer: C

Explanation:

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

NEW QUESTION 105

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 107

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 109

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

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

Answer: D

Explanation:

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

NEW QUESTION 114

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 117

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

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

Answer: C

Explanation:

In AWS Data Pipeline, you can define Amazon SNS alarms to trigger on actMties such as success, failure, or delay by creating an alarm object and referencing it in the onFail, onSuccess, or onLate slots of the actMty object.

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

NEW QUESTION 119

Mike is appointed as Cloud Consultant in ExamKiller.com. ExamKiller has the following VPCs set-up in the US East Region:

A VPC with CIDR block 10.10.0.0/16, a subnet in that VPC with CIDR block 10.10.1.0/24 A VPC with CIDR block 10.40.0.0/16, a subnet in that VPC with CIDR block 10.40.1.0/24

ExamKiller.com is trying to establish network connection between two subnets, a subnet with CIDR block 10.10.1.0/24 and another subnet with CIDR block 10.40.1.0/24. Which one of the following solutions should IVjike recommend to ExamKiller.com?

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

Answer: C

Explanation:

A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IP addresses. EC2 instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, or with a VPC in another AWS account within a single region.

AWS uses the existing infrastructure of a VPC to create a VPC peering connection; it is neither a gateway nor a VPN connection, and does not rely on a separate piece of physical hardware.

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

NEW QUESTION 120

Can Provisioned IOPS be used on RDS instances launched in a VPC?

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

Answer: B

Explanation:

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

Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/RDSFAQ.PIOPS.html>

NEW QUESTION 125

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 127

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 132

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 136

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 141

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

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

Answer: C

Explanation:

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

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

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

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

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

NEW QUESTION 145

You are looking to migrate your Development (Dev) and Test environments to AWS. You have decided to use separate AWS accounts to host each environment. You plan to link each accounts bill to a Master AWS account using Consolidated Billing. To make sure you Keep within budget you would like to implement a way for administrators in the Master account to have access to stop, delete and/or terminate resources in both the Dev and Test accounts. Identify which option will allow you to achieve this goal.

- A. Create IAM users in the Master account with full Admin permission
- B. Create cross-account roles in the Dev and Test accounts that grant the Master account access to the resources in the account by inheriting permissions from the Master account.
- C. Create IAM users and a cross-account role in the Master account that grants full Admin permissions to the Dev and Test accounts.
- D. Create IAM users in the Master account Create cross-account roles in the Dev and Test accounts that have full Admin permissions and grant the Master account access.
- E. Link the accounts using Consolidated Billin
- F. This will give IAM users in the Master account access to resources in the Dev and Test accounts

Answer: C

NEW QUESTION 148

To serve Web traffic for a popular product your chief financial officer and IT director have purchased 10 ml large heavy utilization Reserved Instances (RIs) evenly spread across two availability zones: Route 53 is used to deliver the traffic to an Elastic Load Balancer (ELB). After several months, the product grows even more

popular and you need additional capacity. As a result, your company purchases two C3.2xlarge medium utilization RIs. You register the two c3 2xlarge instances with your ELB and quickly find that the m1 large instances are at 100% of capacity and the c3 2xlarge instances have significant capacity that's unused. Which option is the most cost effective and uses EC2 capacity most effectively?

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

Answer: B

NEW QUESTION 151

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 to 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 CloudWatch.
- 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 153

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

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

Answer: D

NEW QUESTION 156

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 158

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

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

Answer: A

NEW QUESTION 163

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 167

Company B is launching a new game app for mobile devices. Users will log into the game using their existing social media account to streamline data capture. Company B would like to directly save player data and scoring information from the mobile app to a DynamoDB table named Score Data When a user saves their game the progress data will be stored to the Game state S3 bucket. What is the best approach for storing data to DynamoDB and S3?

- A. Use an EC2 Instance that is launched with an EC2 role providing access to the Score Data DynamoDB table and the GameState S3 bucket that communicates with the mobile app via web services.
- B. Use temporary security credentials that assume a role providing access to the Score Data DynamoDB table and the Game State S3 bucket using web identity federation.
- C. Use Login with Amazon allowing users to sign in with an Amazon account providing the mobile app with access to the Score Data DynamoDB table and the Game State S3 bucket.
- D. Use an IAM user with access credentials assigned a role providing access to the Score Data DynamoDB table and the Game State S3 bucket for distribution with the mobile app.

Answer: B

NEW QUESTION 172

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

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

Answer: AC

NEW QUESTION 175

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

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

Answer: BD

NEW QUESTION 178

Refer to the architecture diagram above of a batch processing solution using Simple Queue Service (SQS) to set up a message queue between EC2 instances which are used as batch processors Cloud Watch monitors the number of Job requests (queued messages) and an Auto Scaling group adds or deletes batch servers automatically based on parameters set in Cloud Watch alarms. You can use this architecture to implement which of the following features in a cost effective and efficient manner?

- A. Reduce the overall time for executing jobs through parallel processing by allowing a busy EC2 instance that receives a message to pass it to the next instance in a daisy-chain setup.
- B. Implement fault tolerance against EC2 instance failure since messages would remain in SQS and work can continue with recovery of EC2 instances implement fault tolerance against SQS failure by backing up messages to S3.
- C. Implement message passing between EC2 instances within a batch by exchanging messages through SQS.
- D. Coordinate number of EC2 instances with number of job requests automatically thus Improving cost effectiveness.
- E. Handle high priority jobs before lower priority jobs by assigning a priority metadata field to SQS messages.

Answer: D

NEW QUESTION 180

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

- A. Create an AWS Elastic Beanstalk application using the custom web server platform

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

Answer: AD

NEW QUESTION 183

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

How should they architect their solution?

- A. Route payment requests through two NAT instances setup for High Availability and whitelist the Elastic IP addresses attached to the EC2 instances.
- B. Whitelist the VPC Internet Gateway Public IP and route payment requests through the Internet Gateway.
- C. Whitelist the ELB IP addresses and route payment requests from the Application servers through the ELB.
- D. Automatically assign public IP addresses to the application instances in the Auto Scaling group and run a script on boot that adds each instance's public IP address to the payment validation whitelist API.

Answer: D

NEW QUESTION 185

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

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

Answer: B

NEW QUESTION 186

Which is a valid Amazon Resource name (ARN) for IAM?

- A. aws:iam::123456789012:instance-profile/Nebserver
- B. arn:aws:iam::123456789012:instance-profile/Webserver
- C. 123456789012:aws:iam::instance-profile/Nebserver
- D. arn:aws:iam::123456789012::instance-profile/Nebserver

Answer: B

NEW QUESTION 190

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 194

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 throughpu

Answer: B

NEW QUESTION 198

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

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

Answer: A

NEW QUESTION 199

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 action 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 main 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 201

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

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

Answer: B

NEW QUESTION 205

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

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

Answer: C

NEW QUESTION 210

A customer is deploying an SSL enabled web application to AWS and would like to implement a separation of roles between the EC2 service administrators that

are entitled to login to instances as well as making API calls and the security officers who will maintain and have exclusive access to the application's X.509 certificate that contains the private key.

- A. Upload the certificate on an S3 bucket owned by the security officers and accessible only by EC2 Role of the web servers.
- B. Configure the web servers to retrieve the certificate upon boot from an CloudHSM is managed by the security officers.
- C. Configure system permissions on the web servers to restrict access to the certificate only to the authority security officers
- D. Configure IAM policies authorizing access to the certificate store only to the security officers and terminate SSL on an ELB.

Answer: D

NEW QUESTION 212

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

Which approach provides a cost effective scalable mitigation to this kind of attack?

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

Answer: C

NEW QUESTION 217

You are designing Internet connectMty for your VPC. The Web servers must be available on the Internet. The application must have a highly available architecture. Which alternatives should you consider? (Choose 2 answers)

- A. Configure a NAT instance in your VPC Create a default route via the NAT instance and associate itwith all subnets Configure a DNS A record that points to the NAT instance public IP address.
- B. Configure a CloudFront distribution and configure the origin to point to the private IP addresses of your Web servers Configure a Route53 CNAME record to your CloudFront distribution.
- C. Place all your web servers behind ELB Configure a Route53 CNAME to point to the ELB DNS name.
- D. Assign EIPs to all web server
- E. Configure a Route53 record set with all EIPs, with health checks and DNS failover.
- F. Configure ELB with an EIP Place all your Web servers behind ELB Configure a Route53 A record that points to the EIP.

Answer: CD

NEW QUESTION 219

Your application is using an ELB in front of an Auto Scaling group of web/application servers deployed across two AZs and a Multi-AZ RDS Instance for data persistence.

The database CPU is often above 80% usage and 90% of I/O operations on the database are reads. To improve performance you recently added a single-node Memcached ElastiCache Cluster to cache frequent DB query results. In the next weeks the overall workload is expected to grow by 30%.

Do you need to change anything in the architecture to maintain the high availability or the application with the anticipated additional load? Why?

- A. Yes, you should deploy two Memcached ElastiCache Clusters in different AZs because the RDS instance will not be able to handle the load if the cache node fails.
- B. No, if the cache node fails you can always get the same data from the DB without having any availability impact.
- C. No, if the cache node fails the automated ElastiCache node recovery feature will prevent any availability impact.
- D. Yes, you should deploy the Memcached ElastiCache Cluster with two nodes in the same AZ as the RDS DB master instance to handle the load if one cache node fails.

Answer: A

NEW QUESTION 220

You control access to S3 buckets and objects with:

- A. Identity and Access Management (IAM) Policies.
- B. Access Control Lists (ACLs).
- C. Bucket Policies.
- D. All of the above

Answer: D

NEW QUESTION 224

The following policy can be attached to an IAM group. It lets an IAM user in that group access a "home directory" in AWS S3 that matches their user name using the console.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": ["s3:*"], "Effect": "Allow",
      "Resource": ["arn:aws:s3:::zbucket-name"], "Condition":{"StringLike":{"s3:prefix":["home/${aws:username}/*"]}}
```

```
}!  
{  
  "Action":["s3:*"], "Effect":"Allow",  
  "Resource": ["arn:aws:s3:::bucket-name/home/${aws:username}/*"]  
}  
}
```

- A. True
- B. False

Answer: B

NEW QUESTION 227

You have launched an EC2 instance with four (4) 500 GB EBS Provisioned IOPS volumes attached. The EC2 instance is EBS-Optimized and supports 500 Mbps throughput between EC2 and EBS. The four EBS volumes are configured as a single RAID 0 device, and each Provisioned IOPS volume is provisioned with 4,000 IOPS (4,000 16KB reads or writes), for a total of 16,000 random IOPS on the instance. The EC2 instance initially delivers the expected 16,000 IOPS random read and write performance. Sometime later, in order to increase the total random I/O performance of the instance, you add an additional two 500 GB EBS Provisioned IOPS volumes to the RAID. Each volume is provisioned to 4,000 IOPS like the original four, for a total of 24,000 IOPS on the EC2 instance. Monitoring shows that the EC2 instance CPU utilization increased from 50% to 70%, but the total random IOPS measured at the instance level does not increase at all. What is the problem and a valid solution?

- A. The EBS-Optimized throughput limits the total IOPS that can be utilized; use an EBSOptimized instance that provides larger throughput.
- B. Small block sizes cause performance degradation, limiting the I/O throughput; configure the instance device driver and filesystem to use 64KB blocks to increase throughput.
- C. The standard EBS Instance root volume limits the total IOPS rate; change the instance root volume to also be a 500GB 4,000 Provisioned IOPS volume.
- D. Larger storage volumes support higher Provisioned IOPS rates; increase the provisioned volume storage of each of the 6 EBS volumes to 1TB.
- E. RAID 0 only scales linearly to about 4 devices; use RAID 0 with 4 EBS Provisioned IOPS volumes, but increase each Provisioned IOPS EBS volume to 6,000 IOPS.

Answer: C

NEW QUESTION 228

A newspaper organization has a on-premises application which allows the public to search its back catalogue and retrieve individual newspaper pages via a website written in Java. They have scanned the old newspapers into JPEGs (approx 17TB) and used Optical Character Recognition (OCR) to populate a commercial search product. The hosting platform and software are now end of life and the organization wants to migrate its archive to AWS and produce a cost efficient architecture and still be designed for availability and durability. Which is the most appropriate?

- A. Use S3 with reduced redundancy to store and serve the scanned files, install the commercial search application on EC2 Instances and configure with auto-scaling and an Elastic Load Balancer.
- B. Model the environment using CloudFormation use an EC2 instance running Apache webserver and an open source search application, stripe multiple standard EBS volumes together to store the JPEGs and search index.
- C. Use S3 with standard redundancy to store and serve the scanned files, use CloudSearch for query processing, and use Elastic Beanstalk to host the website across multiple availability zones.
- D. Use a single-AZ RDS MySQL instance to store the search index and the JPEG images use an EC2 instance to serve the website and translate user queries into SQL.
- E. Use a CloudFront download distribution to serve the JPEGs to the end users and install the current commercial search product, along with a Java Container on the website on EC2 instances and use Route53 with DNS round-robin.

Answer: C

NEW QUESTION 230

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

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

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

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