



Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

NEW QUESTION 1

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 2

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

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

Answer: C

Explanation:

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

NEW QUESTION 3

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

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

Answer: B

Explanation:

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

Reference:

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

NEW QUESTION 4

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

- A. The type of customer gateway, such as Cisco ASA, Juniper J-Series, Juniper SSG, Yamaha.
- B. Elastic IP ranges that the organization wants to advertise over the VPN connection to the VPC.
- C. Internet-routable IP address (static) of the customer gateway's external interface.
- D. Border Gateway Protocol (BGP) Autonomous System Number (ASN) of the customer 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 5

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

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

Answer: A

Explanation:

The maximum number of VPCs allowed per region is 5.

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

NEW QUESTION 6

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

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 7

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

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

Answer: D

Explanation:

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

The AWS Direct Connect Location to connect to.

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

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

NEW QUESTION 8

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 9

A user is planning to host a web server as well as an app server on a single EC2 instance which is a part of the public subnet of a VPC. How can the user setup to have two separate public IPs and separate security groups for both the application as well as the web server?

- A. Launch VPC with two separate subnets and make the instance a part of both the subnets.
- B. Launch a VPC instance with two network interface
- C. Assign a separate security group and elastic IP to them.
- D. Launch a VPC instance with two network interface
- E. Assign a separate security group to each and AWS will assign a separate public IP to them.
- F. Launch a VPC with ELB such that it redirects requests to separate VPC instances of the public subne

Answer: B

Explanation:

If you need to host multiple websites(with different IPs) on a single EC2 instance, the following is the suggested method from AWS.

Launch a VPC instance with two network interfaces

Assign elastic IPs from VPC EIP pool to those interfaces (Because, when the user has attached more than one network interface with an instance, AWS cannot

assign public IPs to them.)
Assign separate Security Groups if separate Security Groups are needed
This scenario also helps for operating network appliances, such as firewalls or load balancers that have multiple private IP addresses for each network interface.
Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/MultipleIP.html>

NEW QUESTION 10

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 10

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

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

Answer: C

Explanation:

A user authenticating with Amazon Cognito will go through a multi-step process to bootstrap their credentials. Amazon Cognito has two different flows for authentication with public providers: enhanced and basic.
Reference: <http://docs.aws.amazon.com/cognito/devguide/identity/concepts/authentication-flow/>

NEW QUESTION 11

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

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

Answer: B

Explanation:

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

NEW QUESTION 15

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

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

Answer: C

Explanation:

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

NEW QUESTION 17

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 21

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 24

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 27

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 29

An organization is having an application which can start and stop an EC2 instance as per schedule. The organization needs the MAC address of the instance to be registered with its software. The instance is launched in EC2-CLASSIC. How can the organization update the MAC registration every time an instance is booted?

- A. The organization should write a boot strapping script which will get the MAC address from the instance metadata and use that script to register with the application.
- B. The organization should provide a MAC address as a part of the user data
- C. Thus, whenever the instance is booted the script assigns the fixed MAC address to that instance.
- D. The instance MAC address never change
- E. Thus, it is not required to register the MAC address every time.
- F. AWS never provides a MAC address to an instance; instead the instance ID is used for identifying the instance for any software registration.

Answer: A

Explanation:

AWS provides an on demand, scalable infrastructure. AWS EC2 allows the user to launch On-Demand instances. AWS does not provide a fixed MAC address to the instances launched in EC2-CLASSIC. If the instance is launched as a part of EC2-VPC, it can have an ENI which can have a fixed MAC. However, with EC2-CLASSIC, every time the instance is started or stopped it will have a new MAC address.

To get this MAC, the organization can run a script on boot which can fetch the instance metadata and get the MAC address from that instance metadata. Once the MAC is received, the organization can register that MAC with the software.

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

NEW QUESTION 34

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

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

Answer: B

Explanation:

You can use the action Modify DB Instance, available in the Amazon RDS API, to pass values for the parameters DB Instance Identifier and DB Security Groups specifying the instance ID and the DB Security Groups you want your instance to be part of.

Reference: http://docs.aws.amazon.com/AmazonRDS/latest/APIReference/API_ModifyDBInstance.html

NEW QUESTION 39

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

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

Answer: A

Explanation:

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

NEW QUESTION 41

What types of identities do Amazon Cognito identity pools support?

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

Answer: A

Explanation:

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

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

NEW QUESTION 42

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 44

A user is trying to create a vault in AWS Glacier. The user wants to enable notifications. In which of the below mentioned options can the user enable the notifications from the AWS console?

- A. Glacier does not support the AWS console
- B. Archival Upload Complete
- C. Vault Upload Job Complete
- D. Vault Inventory Retrieval Job Complete

Answer: D

Explanation:

From AWS console the user can configure to have notifications sent to Amazon Simple Notifications Service (SNS). The user can select specific jobs that, on completion, will trigger the notifications such as Vault Inventory Retrieval Job Complete and Archive Retrieval Job Complete.

Reference: <http://docs.aws.amazon.com/amazonglacier/latest/dev/configuring-notifications-console.html>

NEW QUESTION 48

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 49

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/EBSVo|umeTypes.html#EBSVo|umeTypes_pio ps

NEW QUESTION 54

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 59

With Amazon Elastic MapReduce (Amazon EMR) you can analyze and process vast amounts of data. The cluster is managed using an open-source framework

called Hadoop.

You have set up an application to run Hadoop jobs. The application reads data from DynamoDB and generates a temporary file of 100 TBs.

The whole process runs for 30 minutes and the output of the job is stored to S3. Which of the below mentioned options is the most cost effective solution in this case?

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

Answer: B

Explanation:

AWS EC2 Spot Instances allow the user to quote his own price for the EC2 computing capacity. The user can simply bid on the spare Amazon EC2 instances and run them whenever his bid exceeds the current Spot Price. The Spot Instance pricing model complements the On-Demand and Reserved Instance pricing models, providing potentially the most cost-effective option for obtaining compute capacity, depending on the application. The only challenge with a Spot Instance is data persistence as the instance can be terminated whenever the spot price exceeds the bid price.

In the current scenario a Hadoop job is a temporary job and does not run for a longer period. It fetches data from a persistent DynamoDB. Thus, even if the instance gets terminated there will be no data loss and the job can be re-run. As the output files are large temporary files, it will be useful to store data on ephemeral storage for cost savings.

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

NEW QUESTION 64

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

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

Answer: C

Explanation:

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

Instead they should have two separate DBs in each site and setup data replication based on RTO(recovery time objective)of the organization.

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

NEW QUESTION 69

In the context of policies and permissions in AWS IAM, the Condition element is .

- A. crucial while writing the IAM policies
- B. an optional element
- C. always set to null
- D. a mandatory element

Answer: B

Explanation:

The Condition element (or Condition block) lets you specify conditions for when a policy is in effect. The Condition element is optional.

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

NEW QUESTION 71

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

streqi is the short version of the string condition.

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

Answer: A

Explanation:

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

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

NEW QUESTION 76

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

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

Answer: B

Explanation:

You can provision up to 3TB storage and 30,000 IOPS per database instance. For a workload with 50% writes and 50% reads running on a cr1.8xlarge instance, you can realize over 25,000 IOPS for PostgreSQL. However, by provisioning more than this limit, you may be able to achieve lower latency and higher throughput. Your actual realized IOPS may vary from the amount you provisioned based on your database workload, instance type, and database engine choice.

Reference: <https://aws.amazon.com/rds/postgresql/>

NEW QUESTION 77

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 79

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

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

Answer: C

Explanation:

Currently the STS API command GetSessionToken is available to every IAM user in your account without previous permission. In contrast, the GetFederationToken command is restricted and explicit permissions need to be granted so a user can issue calls to this particular Action

Reference: <http://docs.aws.amazon.com/STS/latest/UsingSTS/STSPermission.html>

NEW QUESTION 84

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 tru

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 86

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 87

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

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

D. FALSE

Answer: D

Explanation:

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

NEW QUESTION 89

In Amazon Redshift, how many slices does a dw2.8xlarge node have?

- A. 16
- B. 8
- C. 32
- D. 2

Answer: C

Explanation:

The disk storage for a compute node in Amazon Redshift is divided into a number of slices, equal to the number of processor cores on the node. For example, each DW1.XL compute node has two slices, and each DW2.8XL compute node has 32 slices. Reference: http://docs.aws.amazon.com/redshift/latest/dg/t_Distributing_data.html

NEW QUESTION 91

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

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

Answer: A

Explanation:

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

NEW QUESTION 96

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

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

Answer: C

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. The AWS account provides two platforms: EC2-CLASSIC and EC2-VPC, depending on when the user has created his AWS account and which regions he is using. If the user has created the AWS account after 2013-12-04, it supports only EC2-VPC. In this scenario, since the account is before the required date the supported platform will be EC2-CLASSIC. It is required that the organization creates a VPC as the T2 instances can be launched only as a part of VPC. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/vpc-migrate.html>

NEW QUESTION 98

How does AWS Data Pipeline execute activities on on-premise resources or AWS resources that you manage?

- A. By supplying a Task Runner package that can be installed on your on-premise hosts
- B. None of these
- C. By supplying a Task Runner file that the resources can access for execution
- D. By supplying a Task Runner script that can be installed on your on-premise hosts

Answer: A

Explanation:

To enable running activities using on-premise resources, AWS Data Pipeline does the following: It supplies a Task Runner package that can be installed on your on-premise hosts. This package continuously polls the AWS Data Pipeline service for work to perform.

When it's time to run a particular actMty on your on-premise resources, it will issue the appropriate command to the Task Runner.
Reference: <https://aws.amazon.com/datapipeline/faqs/>

NEW QUESTION 103

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

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

Answer: A

Explanation:

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

NEW QUESTION 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 111

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

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

Answer: A

Explanation:

The two reasons you would need to edit claims issued in a SAML token are: The NameIdentifier claim cannot be the same as the username stored in AD, and The app requires a different set of claim URIs.
Reference:
<https://azure.microsoft.com/en-us/documentation/articles/active-directory-saml-claims-customization/>

NEW QUESTION 114

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

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

Answer: D

Explanation:

AWS CloudHSM provides secure cryptographic key storage to customers by making hardware security modules (HSMs) available in the AWS cloud. AWS CloudHSM requires the following environment before an HSM appliance can be provisioned. A virtual private cloud (VPC) in the region where you want the AWS CloudHSM service.
One private subnet (a subnet with no Internet gateway) in the VPC. The HSM appliance is provisioned into this subnet.
One public subnet (a subnet with an Internet gateway attached). The control instances are attached to this subnet.
An AWS Identity and Access Management (IAM) role that delegates access to your AWS resources to AWS CloudHSM.
An EC2 instance, in the same VPC as the HSM appliance, that has the SafeNet client software installed. This instance is referred to as the control instance and is used to connect to and manage the HSM appliance.
A security group that has port 22 (for SSH) or port 3389 (for RDP) open to your network. This security group is attached to your control instances so you can access them remotely.

NEW QUESTION 119

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 122

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

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

Answer: A

Explanation:

When a request is made, the AWS IAM policy decides whether a given request should be allowed or denied. The evaluation logic follows these rules:

By default, all requests are denied. (In general, requests made using the account credentials for resources in the account are always allowed.)

An explicit allow policy overrides this default.

An explicit deny policy overrides any allows.

In this case since there are explicit deny and explicit allow statements. Thus, the request will be denied since deny overrides allow.

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

NEW QUESTION 127

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

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

Answer: C

Explanation:

In AWS Data Pipeline, you can define Amazon SNS alarms to trigger on actMtys 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 130

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

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

Answer: D

Explanation:

For a single-node configuration in Amazon Redshift, the only option available is the smallest of the two options. The 8XL extra-large nodes are only available in a multi-node configuration

Reference: <http://docs.aws.amazon.com/redshift/latest/mgmt/working-with-clusters.html>

NEW QUESTION 132

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 136

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

- A. Complete the Cross Connect
- B. Configure Redundant Connections with AWS Direct Connect
- C. Create a Virtual Interface
- D. Download Router Configuration

Answer: C

Explanation:

In AWS Direct Connect, your network must support Border Gateway Protocol (BGP) and BGP MD5 authentication, and you need to provide a private Autonomous System Number (ASN) for that to connect to Amazon Virtual Private Cloud (VPC). To connect to public AWS products such as Amazon EC2 and Amazon S3, you will also need to provide a public ASN that you own (preferred) or a private ASN. You have to configure BGP in the Create a Virtual Interface step.

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

NEW QUESTION 140

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 143

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

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

Answer: C

Explanation:

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

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

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

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

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

NEW QUESTION 148

In Amazon ElastiCache, the default cache port is:

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

Answer: D

Explanation:

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

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

NEW QUESTION 150

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 151

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 156

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 ml 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 Cloudwatc
- 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 instance
- 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 161

Your startup wants to implement an order fulfillment process for selling a personalized gadget that needs an average of 3-4 days to produce with some orders taking up to 6 months you expect 10 orders per day on your first day. 1000 orders per day after 6 months and 10,000 orders after 12 months.

Orders coming in are checked for consistency men dispatched to your manufacturing plant for production quality control packaging shipment and payment processing If the product does not meet the quality standards at any stage of the process employees may force the process to repeat a step Customers are notified via email about order status and any critical issues with their orders such as payment failure.

Your case architecture includes AWS Elastic Beanstalk for your website with an RDS MySQL instance for customer data and orders.

How can you implement the order fulfillment process while making sure that the emails are delivered reliably?

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

Answer: C

NEW QUESTION 164

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

- A. Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaimg 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 CloudWate
- 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 169

You are designing a photo-sharing mobile app. The application will store all pictures in a single Amazon S3 bucket. Users will upload pictures from their mobile device directly to Amazon S3 and will be able to view and download their own pictures directly from Amazon S3. You want to configure security to handle potentially millions of users in the most secure manner possible. What should your server-side application do when a new user registers on the photo-sharing mobile application?

- A. Create an IAM use
- B. Update the bucket policy with appropriate permissions for the IAM use
- C. Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.
- D. Create an IAM use
- E. Assign appropriate permissions to the IAM use
- F. Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.
- G. Create a set of long-term credentials using AWS Security Token Service with appropriate permission
- H. Store these credentials in the mobile app and use them to access Amazon S3.
- I. Record the user's information in Amazon RDS and create a role in IAM with appropriate permission
- J. When the user uses their mobile app, create temporary credentials using the AWS Security Token Service "AssumeRole" function
- K. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.
- L. Record the user's information in Amazon DynamoDB
- M. When the user uses their mobile app, create temporary credentials using AWS Security Token Service with appropriate permission
- N. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.

Answer: D

NEW QUESTION 172

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

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

Answer: C

NEW QUESTION 174

You currently operate a web application in the AWS US-East region. The application runs on an auto-scaled layer of EC2 instances and an RDS Multi-AZ database. Your IT security compliance officer has tasked you to develop a reliable and durable logging solution to track changes made to your EC2, IAM, and RDS resources. The solution must ensure the integrity and confidentiality of your log data. Which of these solutions would you recommend?

- A. Create a new CloudTrail trail with one new S3 bucket to store the logs and with the global services option selected. Use IAM roles, S3 bucket policies, and Multi-Factor Authentication (MFA). Delete on the S3 bucket that stores your logs.
- B. Create a new CloudTrail with one new S3 bucket to store the logs. Configure SNS to send log file delivery notifications to your management system. Use IAM roles and S3 bucket policies on the S3 bucket that stores your logs.
- C. Create a new CloudTrail trail with an existing S3 bucket to store the logs and with the global services option selected. Use S3 ACLs and Multi-Factor Authentication (MFA). Delete on the S3 bucket that stores your logs.
- D. Create three new CloudTrail trails with three new S3 buckets to store the logs: one for the AWS Management console, one for AWS SDKs, and one for command-line tools. Use IAM roles and S3 bucket policies on the S3 buckets that store your logs.

Answer: A

NEW QUESTION 175

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 178

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

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

Answer: C

NEW QUESTION 182

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

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

Answer: A

NEW QUESTION 184

You are responsible for a legacy web application whose server environment is approaching end of life. You would like to migrate this application to AWS as quickly as possible, since the application environment currently has the following limitations:

The VM's single 10GB VNIC is almost full; the virtual network interface still uses the 10Gbps driver, which leaves your 100Mbps WAN connection completely underutilized;

It is currently running on a highly customized Windows VM within a VMware environment; You do not have the installation media;

This is a mission-critical application with an RTO (Recovery Time Objective) of 8 hours, RPO (Recovery Point Objective) of 1 hour. How could you best migrate this application to AWS while meeting your business continuity requirements?

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

Answer: A

NEW QUESTION 185

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 187

You must architect the migration of a web application to AWS. The application consists of Linux web servers running a custom web server. You are required to save the logs generated from the application to a durable location.

What options could you select to migrate the application to AWS? (Choose 2)

- A. Create an AWS Elastic Beanstalk application using the custom web server platform.
- B. Specify the web server executable and the application project and source file.
- C. Enable log file rotation to Amazon Simple Storage Service (S3).
- D. Create Dockerfile for the application.
- E. Create an AWS OpsWorks stack consisting of a custom layer.
- F. Create custom recipes to install Docker and to deploy your Docker container using the Dockerfile.
- G. Create custom recipes to install and configure the application to publish the logs to Amazon CloudWatch Logs.
- H. Create Dockerfile for the application.
- I. Create an AWS OpsWorks stack consisting of a Docker layer that uses the Dockerfile.
- J. Create custom recipes to install and configure Amazon Kinesis to publish the logs into Amazon CloudWatch.
- K. Create a Dockerfile for the application.
- L. Create an AWS Elastic Beanstalk application using the Docker platform and the Dockerfile.
- M. Enable logging the Docker configuration to automatically publish the application log.
- N. Enable log file rotation to Amazon S3.

- O. Use VM import/Export to import a virtual machine image of the server into AWS as an AMI
- P. Create an Amazon Elastic Compute Cloud (EC2) instance from AMI, and install and configure the Amazon CloudWatch Logs agent
- Q. Create a new AMI from the instance
- R. Create an AWS Elastic Beanstalk application using the AMI platform and the new AMI.

Answer: AD

NEW QUESTION 192

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 193

Your website is serving on-demand training videos to your workforce. Videos are uploaded monthly in high resolution MP4 format. Your workforce is distributed globally, often on the move, and using company-provided tablets that require the HTTP Live Streaming (HLS) protocol to watch a video. Your company has no video transcoding expertise and it required you may need to pay for a consultant.

How do you implement the most cost-efficient architecture without compromising high availability and quality of video delivery?

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

Answer: C

NEW QUESTION 194

A customer has established an AWS Direct Connect connection to AWS. The link is up and routes are being advertised from the customer's end, however the customer is unable to connect from EC2 instances inside its VPC to servers residing in its datacenter.

Which of the following options provide a viable solution to remedy this situation? (Choose 2 answers)

- A. Add a route to the route table with an IPsec VPN connection as the target.
- B. Enable route propagation to the virtual private gateway (VGW).
- C. Enable route propagation to the customer gateway (CGW).
- D. Modify the route table of all EC2 instances using the 'route' command.
- E. Modify the EC2 instances VPC subnet route table by adding a route back to the customer's on-premises environment.

Answer: AC

NEW QUESTION 199

You are running a news website in the eu-west-1 region that updates every 15 minutes. The website has a world-wide audience; it uses an Auto Scaling group behind an Elastic Load Balancer and an Amazon

RDS database. Static content resides on Amazon S3, and is distributed through Amazon CloudFront. Your Auto Scaling group is set to trigger a scale up event at 60% CPU utilization; you use an Amazon RDS extra large DB instance with 10,000 Provisioned IOPS. Its CPU utilization is around 80%. While freeable memory is in the 2 GB range.

Web analytics reports show that the average load time of your web pages is around 1.5 to 2 seconds, but your SEO consultant wants to bring down the average load time to under 0.5 seconds.

How would you improve page load times for your users? (Choose 3 answers)

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

Answer: ABD

NEW QUESTION 204

You require the ability to analyze a customer's clickstream data on a website so they can do behavioral analysis. Your customer needs to know what sequence of

pages and ads their customer clicked on. This data will be used in real time to modify the page layouts as customers click through the site to increase stickiness and advertising click-through. Which option meets the requirements for captioning and analyzing this data?

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

Answer: B

NEW QUESTION 206

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

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

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

Answer: AB

NEW QUESTION 208

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

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

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

Answer: CDEF

NEW QUESTION 213

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

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

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

Answer: CD

NEW QUESTION 218

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

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

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

Answer: ABD

NEW QUESTION 223

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

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

Answer: C

NEW QUESTION 227

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

Which AWS storage and database architecture meets the requirements of the application?

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

Answer: C

NEW QUESTION 231

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

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

Answer: B

NEW QUESTION 232

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

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

Answer: D

NEW QUESTION 237

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 242

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

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

Answer: ACD

NEW QUESTION 243

You have recently joined a startup company building sensors to measure street noise and air quality in urban areas. The company has been running a pilot deployment of around 100 sensors for 3 months each sensor uploads 1KB of sensor data every minute to a backend hosted on AWS.

During the pilot, you measured a peak of 10 IOPS on the database, and you stored an average of 3GB of sensor data per month in the database.

The current deployment consists of a load-balanced auto scaled Ingestion layer using EC2 instances and a PostgreSQL RDS database with 500GB standard storage.

The pilot is considered a success and your CEO has managed to get the attention of some potential investors. The business plan requires a deployment of at least 100K sensors which needs to be supported by the backend. You also need to store sensor data for at least two years to be able to compare year over year improvements.

To secure funding, you have to make sure that the platform meets these requirements and leaves room for further scaling. Which setup will meet the requirements?

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

Answer: C

NEW QUESTION 244

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 connectivity 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 on-host-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 247

Your team has a tomcat-based Java application you need to deploy into development, test and production environments. After some research, you opt to use Elastic Beanstalk due to its tight integration with your developer tools and RDS due to its ease of management. Your QA team lead points out that you need to roll a sanitized set of production data into your environment on a nightly basis.

Similarly, other software teams in your org want access to that same restored data via their EC2 instances in your VPC. The optimal setup for persistence and security that meets the above requirements would be the following.

- A. Create your RDS instance as part of your Elastic Beanstalk definition and alter its security group to allow access to it from hosts in your application subnets.
- B. Create your RDS instance separately and add its IP address to your application's DB connection strings in your code. Alter its security group to allow access to it from hosts within your VPC's IP address block.
- C. Create your RDS instance separately and pass its DNS name to your app's DB connection string as an environment variable.
- D. Create a security group for client machines and add it as a valid source for DB traffic to the security group of the RDS instance itself.
- E. Create your RDS instance separately and pass its DNS name to your app's DB connection string as an environment variable. Alter its security group to allow access to it from hosts in your application subnets.

Answer: A

NEW QUESTION 250

Your company has an on-premises multi-tier PHP web application, which recently experienced downtime due to a large burst in web traffic due to a company announcement. Over the coming days, you are expecting similar announcements to drive similar unpredictable bursts, and are looking to find ways to quickly improve your infrastructure's ability to handle unexpected increases in traffic.

The application currently consists of 2 tiers: a web tier which consists of a load balancer and several Linux Apache web servers as well as a database tier which hosts a Linux server hosting a MySQL database. Which scenario below will provide full site functionality, while helping to improve the ability of your application in the short timeframe required?

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

I. Create an RDS read replica and setup replication between the RDS instance and on-premises MySQL server to migrate the database.

Answer: C

NEW QUESTION 254

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

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

Answer: C

NEW QUESTION 258

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

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

Answer: D

NEW QUESTION 262

.....

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](#)