

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 organization is planning to host an application on the AWS VPC. The organization wants dedicated instances. However, an AWS consultant advised the organization not to use dedicated instances with VPC as the design has a few limitations. Which of the below mentioned statements is not a limitation of dedicated instances with VPC?

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 4

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 5

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

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

D. aws:CurrentTime

Answer: C

Explanation:

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

NEW QUESTION 6

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

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

Answer: C

Explanation:

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

NEW QUESTION 7

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 8

What bandwidths do AWS Direct Connect currently support?

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

Answer: D

Explanation:

AWS Direct Connection currently supports 1Gbps and 10 Gbps.

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

NEW QUESTION 9

A user is configuring MySQL RDS with PIOPS. What should be the minimum size of DB storage provided by the user?

- A. 1 TB
- B. 50 GB
- C. 5 GB
- D. 100 GB

Answer: D

Explanation:

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

NEW QUESTION 10

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 10

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 12

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

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

Answer: B

Explanation:

Generally AWS EC2 allows running 20 on-demand instances and 100 spot instances at a time. This limit can be increased by requesting at <https://aws.amazon.com/contact-us/ec2-request>. Excluding certain types of instances, the limit is lower than mentioned above. For c4.xlarge, the user can run only 2 on-demand instances at a time.

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

NEW QUESTION 13

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

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

Answer: B

Explanation:

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

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

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

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

NEW QUESTION 17

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 18

What feature of the load balancing service attempts to force subsequent connections to a service to be redirected to the same node as long as it is online?

- A. Node balance
- B. Session retention
- C. Session multiplexing

D. Session persistence

Answer: D

Explanation:

Session persistence is a feature of the load balancing service. It attempts to force subsequent connections to a service to be redirected to the same node as long as it is online.

Reference:

<http://docs.rackspace.com/loadbalancers/api/v1.0/clb-devguide/content/Concepts-d1e233.html>

NEW QUESTION 22

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 26

ExamKiller has three separate departments and each department has their own AWS accounts. The HR department has created a file sharing site where all the on roll employees' data is uploaded. The Admin department uploads data about the employee presence in the office to their DB hosted in the VPC. The Finance department needs to access data from the HR department to know the on roll employees to calculate the salary based on the number of days that an employee is present in the office.

How can ExamKiller setup this scenario?

- A. It is not possible to configure VPC peering since each department has a separate AWS account.
- B. Setup VPC peering for the VPCs of Admin and Finance.
- C. Setup VPC peering for the VPCs of Finance and HR as well as between the VPCs of Finance and Admin.
- D. Setup VPC peering for the VPCs of Admin and HR

Answer: C

Explanation:

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

Reference:

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

NEW QUESTION 31

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 33

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 35

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

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

Answer: C

Explanation:

A Provisioned IOPS (SSD) volume can range in size from 4 GiB to 16 TiB and you can provision up to 20,000 IOPS per volume.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html#EBSVolumeTypes_provisioned_iops

NEW QUESTION 38

Which of the following cache engines does Amazon ElastiCache support?

- A. Amazon ElastiCache supports Memcached and Redis.
- B. Amazon ElastiCache supports Redis and WinCache.
- C. Amazon ElastiCache supports Memcached and Hazelcast.
- D. Amazon ElastiCache supports Memcached only

Answer: A

Explanation:

The cache engines supported by Amazon ElastiCache are Memcached and Redis.
Reference: <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/SelectEngine.html>

NEW QUESTION 43

One of the AWS account owners faced a major challenge in June as his account was hacked and the hacker deleted all the data from his AWS account. This resulted in a major blow to the business.
Which of the below mentioned steps would not have helped in preventing this action?

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

Answer: C

Explanation:

AWS security follows the shared security model where the user is as much responsible as Amazon. If the user wants to have secure access to AWS while hosting applications on EC2, the first security rule to follow is to enable MFA for all users. This will add an added security layer. In the second step, the user should never give his access or secret access keys to anyone as well as store inside programs. The better solution is to use IAM roles. For critical data of the organization, the user should keep an offsite/ in premise backup which will help to recover critical data in case of security breach.

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

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

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

NEW QUESTION 44

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 47

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

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

Answer: B

Explanation:

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

Reference:

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

NEW QUESTION 50

Attempts, one of the three types of items associated with the schedule pipeline in the AWS Data Pipeline, provides robust data management. Which of the following statements is NOT true about Attempts?

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

Answer: C

Explanation:

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

Reference:

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

NEW QUESTION 54

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 58

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 62

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 65

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

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

Answer: A

Explanation:

A service account is a special Google account that can be used by applications to access Google services programmatically. This account belongs to your application or a virtual machine (VM), instead of to an individual end user. Your application uses the service account to call the Google API of a service, so that the users aren't directly involved. A service account can have zero or more pairs of service account keys, which are used to authenticate to Google. A service account key is a public/private keypair generated by Google. Google retains the public key, while the user is given the private key.
Reference: <https://cloud.google.com/iam/docs/service-accounts>

NEW QUESTION 67

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

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

Answer: A

Explanation:

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

NEW QUESTION 70

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

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

Answer: A

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. A VPC peering connection allows the user to route traffic between the peer VPCs using private IP addresses as if they are a part of the same network. This is helpful when one VPC from the same or different AWS account wants to connect with resources of the other VPC. The organization wants to setup that one VPC can connect with all the other VPCs but all other VPCs cannot connect among each other. This can be achieved by configuring VPC peering where one VPC is peered with all the other VPCs, but the other VPCs are not peered to each other. The VPCs are in the same or a separate AWS account and should not have overlapping CIDR blocks.
Reference:
<http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/peering-configurations-full-access.html#many-vpcs-full-access>

NEW QUESTION 74

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 77

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 json script that can be installed on your on-premise hosts

Answer: A

Explanation:

To enable running actMties using on-premise resources, AWS Data Pipeline does the following: It supply 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 78

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 81

You are setting up some EBS volumes for a customer who has requested a setup which includes a RAID (redundant array of inexpensive disks). AWS has some recommendations for RAID setups. Which RAID setup is not recommended for Amazon EBS?

- A. RAID 1 only
- B. RAID 5 only
- C. RAID 5 and RAID 6
- D. RAID 0 only

Answer: C

Explanation:

With Amazon EBS, you can use any of the standard RAID configurations that you can use with a traditional bare metal server, as long as that particular RAID configuration is supported by the operating

system for your instance. This is because all RAID is accomplished at the software level. For greater I/O performance than you can achieve with a single volume, RAID 0 can stripe multiple volumes together; for on-instance redundancy, RAID 1 can mirror two volumes together.

RAID 5 and RAID 6 are not recommended for Amazon EBS because the parity write operations of these RAID modes consume some of the IOPS available to your volumes.

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

NEW QUESTION 83

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 87

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

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

Answer: C

Explanation:

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

NEW QUESTION 90

Who is responsible for modifying the routing tables and networking ACLs in a VPC to ensure that a DB instance is reachable from other instances in the VPC?

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

Answer: B

Explanation:

You are in charge of configuring the routing tables of your VPC as well as the network ACLs rules needed to make your DB instances accessible from all the instances of your VPC that need to communicate with it.

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

NEW QUESTION 94

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

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

Answer: A

Explanation:

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

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

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

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

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

NEW QUESTION 95

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_pio ps

NEW QUESTION 96

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

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

Answer: A

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. The VPC is always specific to a region. The user can create a VPC which can span multiple Availability Zones by adding one or more subnets in each Availability Zone. Each subnet must reside entirely within one Availability Zone and cannot span across zones.

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

NEW QUESTION 97

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

NEW QUESTION 98

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

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

Answer: D

Explanation:

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

NEW QUESTION 102

You're trying to delete an SSL certificate from the IAM certificate store, and you're getting the message "Certificate: <certificate-id> is being used by CloudFront." Which of the following statements is probably the reason why you are getting this error?

- A. Before you can delete an SSL certificate you need to set up https on your server.
- B. Before you can delete an SSL certificate, you need to set up the appropriate access level in IAM
- C. Before you can delete an SSL certificate, you need to either rotate SSL certificates or revert from using a custom SSL certificate to using the default CloudFront certificate.
- D. You can't delete SSL certificates. You need to request it from AWS

Answer: C

Explanation:

CloudFront is a web service that speeds up distribution of your static and dynamic web content, for example, .html, .css, .php, and image files, to end users. Every CloudFront web distribution must be associated either with the default CloudFront certificate or with a custom SSL certificate. Before you can delete an SSL certificate, you need to either rotate SSL certificates (replace the current custom SSL certificate with another custom SSL certificate) or revert from using a custom SSL certificate to using the default CloudFront certificate.
Reference: <http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Troubleshooting.html>

NEW QUESTION 105

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

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

Answer: A

Explanation:

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

NEW QUESTION 106

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

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

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

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

Answer: C

Explanation:

A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IP addresses. EC2 instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, or with a VPC in another AWS account within a single region. AWS uses the existing infrastructure of a VPC to create a VPC peering connection; it is neither a gateway nor a VPN connection, and does not rely on a separate piece of physical hardware.
Reference: <http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-peering.html>

NEW QUESTION 108

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 110

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 115

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

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

Answer: B

Explanation:

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

NEW QUESTION 120

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

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

Answer: D

Explanation:

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

NEW QUESTION 123

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

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

Answer: C

NEW QUESTION 125

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

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

Answer: ADF

NEW QUESTION 129

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 C|oudTrail| 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 mat 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 134

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

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

Answer: C

NEW QUESTION 138

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

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

Answer: AC

NEW QUESTION 142

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

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

Answer: B

NEW QUESTION 147

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 148

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 153

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 HLS
- 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 HLS
- 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 156

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

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

E. And verify network traffic is leveraging the DirectConnect connection.

Answer: D

NEW QUESTION 158

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

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

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

Answer: CE

NEW QUESTION 163

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

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

Answer: B

NEW QUESTION 166

Dave is the main administrator in Example Corp., and he decides to use paths to help delineate the users in the company and set up a separate administrator group for each path-based division. Following is a subset of the full list of paths he plans to use:

. /marketing
. /sales
. Legal

Dave creates an administrator group for the marketing part of the company and calls it MktAdmin. He assigns it the /marketing path. The group's ARN is arn:aws:iam::123456789012:group/marketing/MktAdmin.

Dave assigns the following policy to the MktAdmin group that gives the group permission to use all IAM actions with all groups and users in the /marketing path. The policy also gives the MktAdmin group permission to perform any AWS S3 actions on the objects in the portion of the corporate bucket.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "iam:*",
      "Resource": [
        "arn:aws:iam::123456789012:group/marketing/*",
        "arn:aws:iam::123456789012:user/marketing/*"
      ]
    },
    {
      "Effect": "Allow",
      "Action": "s3:*",
      "Resource": "arn:aws:s3:::example_bucket/marketing/*"
    },
    {
      "Effect": "Allow",
      "Action": "s3:ListBucket",
      "Resource": "arn:aws:s3:::example_bucket",
      "Condition": {
        "StringLike": {
          "s3:prefix": "marketing/"
        }
      }
    }
  ]
}
```

- A. True
- B. False

Answer: B

NEW QUESTION 170

A company is running a batch analysis every hour on their main transactional DB, running on an RDS MySQL instance, to populate their central Data Warehouse running on Redshift. During the execution of the batch, their transactional applications are very slow. When the batch completes they need to update the top management dashboard with the new data. The dashboard is produced by another system running on-premises that is currently started when a manually-sent email notifies that an update is required. The on-premises system cannot be modified because it is managed by another team.

How would you optimize this scenario to solve performance issues and automate the process as much as possible?

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

Answer: A

NEW QUESTION 172

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

Which of the following options would you consider?

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

Answer: A

NEW QUESTION 177

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

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

Answer: B

NEW QUESTION 182

Your system recently experienced down time during the troubleshooting process. You found that a new administrator mistakenly terminated several production EC2 instances.

Which of the following strategies will help prevent a similar situation in the future? The administrator still must be able to: launch, start stop, and terminate development resources. launch and start production instances.

- A. Create an IAM user, which is not allowed to terminate instances by leveraging production EC2 termination protection.
- B. Leverage resource based tagging, along with an IAM user which can prevent specific users from terminating production, EC2 resources.
- C. Leverage EC2 termination protection and multi-factor authentication, which together require users to authenticate before terminating EC2 instances
- D. Create an IAM user and apply an IAM role which prevents users from terminating production EC2 instances.

Answer: B

NEW QUESTION 186

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 191

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

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

Answer: D

NEW QUESTION 195

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 Nlain Route Table with the new EIP
- C. Create 15 Security Group rules to block the attacking IP addresses over port 80
- D. Create an inbound NACL (Network Access control list) associated with the web tier subnet with deny rules to block the attacking IP addresses

Answer: D

NEW QUESTION 197

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

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

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

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

Answer: D

NEW QUESTION 199

A web-startup runs its very successful social news application on Amazon EC2 with an Elastic Load Balancer, an Auto-Scaling group of Java/Tomcat application-servers, and DynamoDB as data store. The main web-application best runs on m2 x large instances since it is highly memory- bound Each new deployment requires semi-automated creation and testing of a new AMI for the application servers which takes quite a while ana is therefore only done once per week.

Recently, a new chat feature has been implemented in nodejs and wails to be integrated in the architecture. First tests show that the new component is CPU bound Because the company has some experience with using Chef, they decided to streamline the deployment process and use AWS Ops Works as an application life cycle tool to simplify management of the application and reduce the deployment cycles.

What configuration in AWS Ops Works is necessary to integrate the new chat module in the most cost-efficient and filexible way?

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

Answer: C

NEW QUESTION 203

How can an EBS volume that is currently attached to an EC2 instance be migrated from one Availability Zone to another?

- A. Detach the volume and attach it to another EC2 instance in the other AZ.
- B. Simply create a new volume in the other AZ and specify the original volume as the source.
- C. Create a snapshot of the volume, and create a new volume from the snapshot in the other AZ.
- D. Detach the volume, then use the ec2-migrate-volume command to move it to another AZ.

Answer: C

NEW QUESTION 205

After launching an instance that you intend to serve as a NAT (Network Address Translation) device in a public subnet you modify your route tables to have the NAT device be the target of internet bound traffic of your private subnet. When you try and make an outbound connection to the internet from an instance in the private subnet, you are not successful. Which of the following steps could resolve the issue?

- A. Disabling the Source/Destination Check attribute on the NAT instance
- B. Attaching an Elastic IP address to the instance in the private subnet
- C. Attaching a second Elastic Network Interface (ENI) to the NAT instance, and placing it in the private subnet
- D. Attaching a second Elastic Network Interface (ENI) to the instance in the private subnet, and placing it in the public subnet

Answer: A

NEW QUESTION 209

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

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

Answer: ACDF

NEW QUESTION 214

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

- A. Implement third party volume encryption tools
- B. Implement SSL/TLS for all services running on the sewer

- C. Encrypt data inside your applications before storing it on EBS
- D. Encrypt data using native data encryption drivers at the file system level
- E. Do nothing as EBS volumes are encrypted by default

Answer: ACD

NEW QUESTION 215

Your firm has uploaded a large amount of aerial image data to S3. In the past, in your on-premises environment, you used a dedicated group of servers to process this data and used Rabbit MQ - An open source messaging system to get job information to the servers. Once processed the data would go to tape and be shipped offsite. Your manager told you to stay with the current design, and leverage AWS archival storage and messaging services to minimize cost. Which is correct?

- A. Use SQS for passing job messages use Cloud Watch alarms to terminate EC2 worker instances when they become idle
- B. Once data is processed, change the storage class of the S3 objects to Reduced Redundancy Storage.
- C. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SQS. Once data is processed,
- D. Change the storage class of the S3 objects to Reduced Redundancy Storage
- E. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SQS. Once data is processed, change the storage class of the S3 objects to Glacier.
- F. Use SNS to pass job messages use Cloud Watch alarms to terminate spot worker instances when they become idle
- G. Once data is processed, change the storage class of the S3 object to Glacier.

Answer: D

NEW QUESTION 217

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 221

Your company is in the process of developing a next generation pet collar that collects biometric information to assist families with promoting healthy lifestyles for their pets. Each collar will push 30kb of biometric data in JSON format every 2 seconds to a collection platform that will process and analyze the data providing health trending information back to the pet owners and veterinarians via a web portal. Management has tasked you to architect the collection platform ensuring the following requirements are met.

Provide the ability for real-time analytics of the inbound biometric data. Ensure processing of the biometric data is highly durable, elastic and parallel. The results of the analytic processing should be persisted for data mining.

Which architecture outlined below will meet the initial requirements for the collection platform?

- A. Utilize S3 to collect the inbound sensor data, analyze the data from S3 with a daily scheduled Data Pipeline and save the results to a Redshift Cluster.
- B. Utilize Amazon Kinesis to collect the inbound sensor data, analyze the data with Kinesis clients and save the results to a Redshift cluster using EMR.
- C. Utilize SQS to collect the inbound sensor data, analyze the data from SQS with Amazon Kinesis and save the results to a Microsoft SQL Server RDS instance.
- D. Utilize EMR to collect the inbound sensor data, analyze the data from EMR with Amazon Kinesis and save the results to DynamoDB.

Answer: B

NEW QUESTION 225

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 227

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

The database CPU is often above 80% usage and 90% of I/O operations on the database are reads. To improve performance you recently added a single-node

Memcached ElastiCache Cluster to cache frequent DB query results. In the next weeks the overall workload is expected to grow by 30%. Do you need to change anything in the architecture to maintain the high availability or the application with the anticipated additional load? Why?

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

Answer: A

NEW QUESTION 229

You control access to S3 buckets and objects with:

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

Answer: D

NEW QUESTION 234

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

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

Answer: C

NEW QUESTION 236

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

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

Answer: C

NEW QUESTION 240

.....

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](#)