



**Amazon**

## **Exam Questions AWS-Certified-Solutions-Architect-Professional**

Amazon AWS Certified Solutions Architect Professional

#### NEW QUESTION 1

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

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

**Answer:** D

#### Explanation:

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

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

#### NEW QUESTION 2

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

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

**Answer:** B

#### Explanation:

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

Reference:

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

#### NEW QUESTION 3

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

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

**Answer:** B

#### Explanation:

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

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

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

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

[http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC\\_VPN.html](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html)

#### NEW QUESTION 4

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

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

**Answer:** B

#### Explanation:

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

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

#### NEW QUESTION 5

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

instances with VPC?

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

**Answer: C**

**Explanation:**

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

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

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

**NEW QUESTION 6**

In 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 7**

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

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 8**

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 9**

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 10

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

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

**Answer: D**

#### Explanation:

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

Reference: [http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage\\_ElementDescriptions.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html)

#### NEW QUESTION 10

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

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

**Answer: A**

#### Explanation:

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

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

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

Reference: [http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/USVPC\\_creating\\_basic\\_lb.html](http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/USVPC_creating_basic_lb.html)

#### NEW QUESTION 14

In Amazon ElastiCache, the failure of a single cache node can have an impact on the availability of your application and the load on your back-end database while ElastiCache provisions a replacement for the failed cache node and it get repopulated. Which of the following is a solution to reduce this potential availability impact?

- A. Spread your memory and compute capacity over fewer number of cache nodes, each with smaller capacity.
- B. Spread your memory and compute capacity over a larger number of cache nodes, each with smaller capacity.
- C. Include fewer number of high capacity nodes.
- D. Include a larger number of cache nodes, each with high capacity

**Answer: B**

#### Explanation:

In Amazon ElastiCache, the number of cache nodes in the cluster is a key factor in the availability of your cluster running Memcached. The failure of a single cache node can have an impact on the availability of your application and the load on your back-end database while ElastiCache provisions a replacement for the failed cache node and it get repopulated. You can reduce this potential availability impact by spreading your memory and compute capacity over a larger number of cache nodes, each with smaller capacity, rather than using a fewer number of high capacity nodes.

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

#### NEW QUESTION 16

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

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

**Answer: A**

#### Explanation:

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

#### NEW QUESTION 20

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](http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.html)

**NEW QUESTION 23**

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](http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html)

**NEW QUESTION 24**

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

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

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

**Answer:** A

**Explanation:**

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

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

Reference:

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

**NEW QUESTION 29**

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

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

**Answer:** D

**Explanation:**

The two policies that you attach to an IAM role are the access policy and the trust policy.

Remember that adding an account to the trust policy of a role is only half of establishing the trust relationship. By default, no users in the trusted accounts can assume the role until the administrator for that account grants the users the permission to assume the role by adding the Amazon Resource Name (ARN) of the role to an Allow element for the sts:AssumeRole action.

Reference: [http://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles\\_manage\\_modify.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_manage_modify.html)

**NEW QUESTION 31**

One of your AWS Data Pipeline activities has failed consequently and has entered a hard failure state after retrying thrice. You want to try it again. Is it possible to

increase the number of automatic retries to more than thrice?

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

**Answer:** D

**Explanation:**

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

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

**NEW QUESTION 32**

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 tunin

**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 35**

How many cg1.4xlarge on-demand instances can a user run in one region without taking any limit increase approval from AWS?

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

**Answer:** B

**Explanation:**

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

Reference: [http://docs.aws.amazon.com/general/latest/gr/aws\\_service\\_limits.html#limits\\_ec2](http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_ec2)

**NEW QUESTION 40**

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

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

**Answer:** A

**Explanation:**

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

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

**NEW QUESTION 45**

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

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

**Answer:** C

**Explanation:**

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

replication and Multi-AZ which can be used to achieve cross AZ redundancy.  
Reference: <https://aws.amazon.com/elasticache/>

#### NEW QUESTION 46

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 dat
- 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 49

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 ligh

**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](http://d36cz9buwru1tt.cloudfront.net/AWS_Disaster_Recovery.pdf)

#### NEW QUESTION 53

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 identitie

**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-pool|s/>

#### NEW QUESTION 54

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

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

**Answer:** A

#### Explanation:

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

[http://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_credentials\\_temp\\_control-access\\_disable-perms.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp_control-access_disable-perms.html)

#### NEW QUESTION 58

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](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Appendix_Limits.html)

#### NEW QUESTION 60

An organization has created 5 IAM users. The organization wants to give them the same login ID but different passwords. How can the organization achieve this?

- A. The organization should create each user in a separate region so that they have their own URL to login
- B. The organization should create a separate login ID but give the IAM users the same alias so that each one can login with their alias
- C. It is not possible to have the same login ID for multiple IAM users of the same account
- D. The organization should create various groups and add each user with the same login ID to different group
- E. The user can login with their own group ID

**Answer:** C

#### Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. Whenever the organization is creating an IAM user, there should be a unique ID for each user. It is not possible to have the same login ID for multiple users. The names of users, groups, roles, instance profiles must be alphanumeric, including the following common characters: plus (+), equal (=), comma (,), period (.), at (@), and dash (-).

Reference: [http://docs.aws.amazon.com/IAM/latest/UserGuide/Using\\_SettingUpUser.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_SettingUpUser.html)

#### NEW QUESTION 61

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 62

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

- A. No, since the Provisioned IOPS 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 Provisioned IOPS 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](http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html#EBSVolumeTypes_provisioned_iops)

#### NEW QUESTION 64

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 66

You have been given the task to define multiple AWS Data Pipeline schedules for different activities in the same pipeline. Which of the following would successfully accomplish this task?

- A. Creating multiple pipeline definition files
- B. Defining multiple pipeline definitions in your schedule objects file and associating the desired schedule to the correct actMty via its schedule field
- C. Defining multiple schedule objects in your pipeline definition file and associating the desired schedule to the correct actMty via its schedule field
- D. Defining multiple schedule objects in the schedule field

**Answer:** C

**Explanation:**

To define multiple schedules for different actMties in the same pipeline, in AWS Data Pipeline, you should define multiple schedule objects in your pipeline definition file and associate the desired schedule to the correct actMty via its schedule field. As an example of this, it could allow you to define a pipeline in which log files are stored in Amazon S3 each hour to drive generation of an aggregate report once a day. Reference: <https://aws.amazon.com/datapipeline/faqs/>

**NEW QUESTION 68**

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 70**

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 73**

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](http://media.amazonwebservices.com/pdf/AWS_Security_Whitepaper.pdf)

**NEW QUESTION 78**

An organization is making software for the CIA in US

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 80**

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 83**

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

**NEW QUESTION 87**

Which of the following is true while using an IAM role to grant permissions to applications running on Amazon EC2 instances?

- A. All applications on the instance share the same role, but different permissions.
- B. All applications on the instance share multiple roles and permissions.
- C. Multiple roles are assigned to an EC2 instance at a time.
- D. Only one role can be assigned to an EC2 instance at a time

**Answer:** D

**Explanation:**

Only one role can be assigned to an EC2 instance at a time, and all applications on the instance share the same role and permissions. Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/role-usecase-ec2app.html>

**NEW QUESTION 91**

When using string conditions within IAM, short versions of the available comparators can be used instead of the more verbose ones. streq 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, `StringEqualsIgnoreCase` 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 93**

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 95**

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

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

**Answer:** C

**Explanation:**

AWS Direct Connect itself has no specific resources for you to control access to. Therefore, there are no AWS Direct Connect ARNs for you to use in an IAM policy. You use an asterisk (\*) as the resource when writing a policy to control access to AWS Direct Connect actions.

Reference: [http://docs.aws.amazon.com/directconnect/latest/UserGuide/using\\_iam.html](http://docs.aws.amazon.com/directconnect/latest/UserGuide/using_iam.html)

**NEW QUESTION 98**

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

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

**Answer:** B

**Explanation:**

Regardless of how your Lambda function is invoked, AWS Lambda always executes the function. At the time you create a Lambda function, you specify an IAM role that AWS Lambda can assume to execute your Lambda function on your behalf. This role is also referred to as the execution role.

Reference: <http://docs.aws.amazon.com/lambda/latest/dg/lambda-dg.pdf>

**NEW QUESTION 100**

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 104**

An organization is having a VPC for the HR department, and another VPC for the Admin department. The HR department requires access to all the instances running in the Admin VPC while the Admin department requires access to all the resources in the HR department. How can the organization setup this scenario?

- A. Setup VPC peering between the VPCs of Admin and HR.
- B. Setup ACL with both VPCs which will allow traffic from the CIDR of the other VPC.
- C. Setup the security group with each VPC which allows traffic from the CIDR of another VPC.
- D. It is not possible to connect resources of one VPC from another 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.

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

**NEW QUESTION 107**

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

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

**Answer:** A

**Explanation:**

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

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

Reference:

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

**NEW QUESTION 108**

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 dMded 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](http://docs.aws.amazon.com/redshift/latest/dg/t_Distributing_data.html)

**NEW QUESTION 111**

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

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

**Answer:** C

**Explanation:**

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

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

**NEW QUESTION 112**

Out of the striping options available for the EBS volumes, which one has the following disadvantage: 'Doubles the amount of I/O required from the instance to EBS compared to RAID 0, because you're mirroring all writes to a pair of volumes, limiting how much you can stripe.'?

- A. Raid 1
- B. Raid 0
- C. RAID 1+0 (RAID 10)
- D. Raid 2

**Answer:** C

**Explanation:**

RAID 1+0 (RAID 10) doubles the amount of I/O required from the instance to EBS compared to RAID 0, because you're mirroring all writes to a pair of volumes, limiting how much you can stripe.

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

**NEW QUESTION 117**

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 119**

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 120**

In Amazon Cognito what is a silent push notification?

- A. It is a push message that is received by your application on a user's device that will not be seen by the user.
- B. It is a push message that is received by your application on a user's device that will return the user's geolocation.
- C. It is a push message that is received by your application on a user's device that will not be heard by the user.
- D. It is a push message that is received by your application on a user's device that will return the user's authentication credentials.

**Answer:** A

**Explanation:**

Amazon Cognito uses the Amazon Simple Notification Service (SNS) to send silent push notifications to devices. A silent push notification is a push message that is received by your application on a user's device that will not be seen by the user.

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

**NEW QUESTION 123**

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 126**

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

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

**Answer:** D

**Explanation:**

The NotAction element lets you specify an exception to a list of actions. Reference:  
[http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage\\_ElementDescriptions.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html)

**NEW QUESTION 127**

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 128**

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 129**

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 134**

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

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

**Answer:** A

**Explanation:**

Amazon ElastiCache cache security groups are only applicable to cache clusters that are not running in an Amazon Virtual Private Cloud environment (VPC). If you are running in an Amazon Virtual Private Cloud, Cache Security Groups is not available in the console navigation pane.  
Reference: <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/CacheSecurityGroup.html>

#### NEW QUESTION 135

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](http://docs.aws.amazon.com/datapipeline/latest/APIReference/API_PollForTask.html)

#### NEW QUESTION 137

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](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html#VPCSubnet)

#### NEW QUESTION 138

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 142

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

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

**Answer: C**

#### Explanation:

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

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

#### NEW QUESTION 146

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 150**

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](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html)

**NEW QUESTION 153**

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.htm>

**NEW QUESTION 158**

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

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

**Answer:** C

**Explanation:**

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

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

**NEW QUESTION 161**

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

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

**Answer:** A

**Explanation:**

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

Reference: [http://docs.aws.amazon.com/IAM/latest/UserGuide/reference\\_policies\\_elements.html#Sid](http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html#Sid)

**NEW QUESTION 164**

You have deployed a web application targeting a global audience across multiple AWS Regions under the domain name.example.com. You decide to use Route53

Latency-Based Routing to serve web requests to users from the region closest to the user. To provide business continuity in the event of server downtime you configure weighted record sets associated with two web servers in separate Availability Zones per region. During a DR test you notice that when you disable all web servers in one of the regions Route53 does not automatically direct all users to the other region. What could be happening? (Choose 2 answers)

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

**Answer: BE**

#### NEW QUESTION 168

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 then 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 servers and re-use the RDS 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 173

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

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

**Answer: A**

#### NEW QUESTION 176

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 177

You are the new IT architect in a company that operates a mobile sleep tracking application.

When activated at night, the mobile app is sending collected data points of 1 kilobyte every 5 minutes to your backend.

The backend takes care of authenticating the user and writing the data points into an Amazon DynamoDB table.

Every morning, you scan the table to extract and aggregate last night's data on a per user basis, and store the results in Amazon S3. Users are notified via Amazon SNS mobile push notifications that new data is available, which is parsed and visualized by the mobile app.

Currently you have around 100k users who are mostly based out of North America. You have been tasked to optimize the architecture of the backend system to lower cost. What would you recommend? Choose 2 answers

- A. Have the mobile app access Amazon DynamoDB directly Instead of JSON files stored on Amazon S3.

- B. Write data directly into an Amazon Redshift cluster replacing both Amazon DynamoDB and Amazon S3.
- C. Introduce an Amazon SQS queue to buffer writes to the Amazon DynamoDB table and reduce provisioned write throughput.
- D. Introduce Amazon ElastiCache to cache reads from the Amazon DynamoDB table and reduce provisioned read throughput.
- E. Create a new Amazon DynamoDB table each day and drop the one for the previous day after its data is on Amazon S3.

**Answer:** AD

#### NEW QUESTION 181

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

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

**Answer:** A

#### NEW QUESTION 186

Your company currently has a 2-tier web application running in an on-premises data center. You have experienced several infrastructure failures in the past two months resulting in significant financial losses. Your CIO is strongly agreeing to move the application to AWS. While working on achieving buy-in from the other company executives, he asks you to develop a disaster recovery plan to help improve Business continuity in the short term. He specifies a target Recovery Time Objective (RTO) of 4 hours and a Recovery Point Objective (RPO) of 1 hour or less. He also asks you to implement the solution within 2 weeks. Your database is 200GB in size and you have a 20Mbps Internet connection. How would you do this while minimizing costs?

- A. Create an EBS backed private AMI which includes a fresh install of your application
- B. Develop a Cloud Formation template which includes your AMI and the required EC2, AutoScaling, and ELB resources to support deploying the application across Multiple- Availability-Zone
- C. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
- D. Deploy your application on EC2 instances within an Auto Scaling group across multiple availability zones
- E. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
- F. Create an EBS backed private AMI which includes a fresh install of your application
- G. Setup a script in your data center to backup the local database every 1 hour and to encrypt and copy the resulting file to an S3 bucket using multi-part upload.
- H. Install your application on a compute-optimized EC2 instance capable of supporting the application's average load
- I. Synchronously replicate transactions from your on-premises database to a database instance in AWS across a secure Direct Connect connection.

**Answer:** A

#### NEW QUESTION 187

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 192

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 196

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 199

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

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

**Answer: C**

#### NEW QUESTION 202

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 204

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 207

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

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

Which design would you choose to meet these requirements?

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

**Answer: A**

#### NEW QUESTION 208

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 209

A corporate web application is deployed within an Amazon Virtual Private Cloud (VPC) and is connected to the corporate data center via an IPsec VPN. The application must authenticate against the on-premises LDAP server. After authentication, each logged-in user can only access an Amazon Simple Storage Space (S3) keyspace specific to that user. Which two approaches can satisfy these objectives? (Choose 2 answers)

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

**Answer: BC**

#### NEW QUESTION 214

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

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

**Answer: C**

#### NEW QUESTION 215

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

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

**Answer: B**

#### NEW QUESTION 216

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 218

Your company produces customer commissioned one-of-a-kind skiing helmets combining high fashion with custom technical enhancements. Customers can show off their individuality on the ski slopes and have access to head-up-displays, GPS rear-view cams and any other technical innovation they wish to embed in the helmet.

The current manufacturing process is data rich and complex including assessments to ensure that the custom electronics and materials used to assemble the helmets are to the highest standards. Assessments are a mixture of human and automated assessments; you need to add a new set of assessment to model the failure modes of the custom electronics using GPUs with CUDA, across a cluster of servers with low latency networking.

What architecture would allow you to automate the existing process using a hybrid approach and ensure that the architecture can support the evolution of processes over time?

- A. Use AWS Data Pipeline to manage movement of data & meta-data and assessments. Use an auto-scaling group of G2 instances in a placement group.
- B. Use Amazon Simple Workflow (SWF) to manage assessments, movement of data & meta-data. Use an auto-scaling group of G2 instances in a placement group.
- C. Use Amazon Simple Workflow (SWF) to manage assessments, movement of data & meta-data. Use an auto-scaling group of C3 instances with SR-IOV (Single Root I/O Virtualization).
- D. Use AWS data Pipeline to manage movement of data & meta-data and assessments. Use an auto-scaling group of C3 with SR-IOV (Single Root I/O virtualization).

**Answer:** B

#### NEW QUESTION 220

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 224

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 225

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

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

**Answer:** A

#### NEW QUESTION 227

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 229

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 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.
- C. Create an IAM role for EC2 that allows list access to objects in the S3 bucket
- D. Launch the instance with the role, and retrieve the role's credentials from the EC2 Instance metadata
- E. Create an IAM user for the application with permissions that allow list access to the S3 bucket
- F. The application retrieves the IAM user credentials from a temporary directory with permissions that allow read access only to the application user.

**Answer: C**

#### NEW QUESTION 234

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

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

**Answer: B**

#### NEW QUESTION 236

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 files 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 Map 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 238

Your application provides data transformation services. Files containing data to be transformed are first uploaded to Amazon S3 and then transformed by a fleet of spot EC2 instances. Files submitted by your premium customers must be transformed with the highest priority. How should you implement such a system?

- A. Use a DynamoDB table with an attribute defining the priority level
- B. Transformation instances will scan the table for tasks, sorting the results by priority level.
- C. Use Route 53 latency based-routing to send high priority tasks to the closest transformation instances.
- D. Use two SQS queues, one for high priority messages, the other for default priority
- E. Transformation instances first poll the high priority queue; if there is no message, they poll the default priority queue.
- F. Use a single SQS queue
- G. Each message contains the priority level
- H. Transformation instances poll high-priority messages first.

**Answer: C**

#### NEW QUESTION 243

When you put objects in Amazon S3, what is the indication that an object was successfully stored?

- A. A HTTP 200 result code and MD5 checksum, taken together, indicate that the operation was successful.
- B. Amazon S3 is engineered for 99.999999999% durability
- C. Therefore there is no need to confirm that data was inserted.
- D. A success code is inserted into the S3 object metadata.
- E. Each S3 account has a special bucket named `_s3_log`
- F. Success codes are written to this bucket with a timestamp and checksum.

**Answer:** A

#### NEW QUESTION 246

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 247

A customer is deploying an SSL enabled web application to AWS and would like to implement a separation of roles between the EC2 service administrators that are entitled to login to instances as well as making API calls and the security officers who will maintain and have exclusive access to the application's X.509 certificate that contains the private key.

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

**Answer:** D

#### NEW QUESTION 248

A web company is looking to implement an intrusion detection and prevention system into their deployed VPC. This platform should have the ability to scale to thousands of instances running inside of the VPC. How should they architect their solution to achieve these goals?

- A. Configure an instance with monitoring software and the elastic network interface (ENI) set to promiscuous mode packet sniffing to see all traffic across the VPC.
- B. Create a second VPC and route all traffic from the primary application VPC through the second VPC where the scalable virtualized IDS/IPS platform resides.
- C. Configure routers running in the VPC using the host-based 'route' commands to send all traffic through the platform to a scalable virtualized IDS/IPS.
- D. Configure each host with an agent that collects all network traffic and sends that traffic to the IDS/IPS platform for inspection.

**Answer:** C

#### NEW QUESTION 249

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, change the storage class of the S3 objects to Glacier.
- 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 254

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 256

You are designing the network infrastructure for an application server in Amazon VPC. Users will access all application instances from the Internet, as well as from an on-premises network. The on-premises network is connected to your VPC over an AWS Direct Connect link. How would you design routing to meet the above requirements?

- A. Configure a single routing table with a default route via the Internet gateway
- B. Propagate a default route via BGP on the AWS Direct Connect customer route
- C. Associate the routing table with all VPC subnets.
- D. Configure a single routing table with a default route via the Internet gateway
- E. Propagate specific routes for the on-premises networks via BGP on the AWS Direct Connect customer route
- F. Associate the routing table with all VPC subnets.
- G. Configure a single routing table with two default routes: one to the Internet via an Internet gateway, the other to the on-premises network via the VPN gateway
- H. Use this routing table across all subnets in the VPC.
- I. Configure two routing tables: one that has a default route via the Internet gateway, and other that has a default route via the VPN gateway
- J. Associate both routing tables with each VPC subnet.

**Answer: A**

#### NEW QUESTION 259

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 262

The AWS IT infrastructure that AWS provides, complies with the following IT security standards, including:

- A. SOC 1/SSAE 16/ISAE 3402 (formerly SAS 70 Type II), SOC 2 and SOC 3
- B. FISMA, DIACAP, and FedRA|V|P
- C. PCI DSS Level 1, ISO 27001, ITAR and FIPS 140-2
- D. HIPAA, Cloud Security Alliance (CSA) and Motion Picture Association of America (NIPAA)
- E. All of the above

**Answer: ABC**

#### NEW QUESTION 264

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

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

- A. True
- B. False

**Answer: B**

#### NEW QUESTION 265

What does elasticity mean to AWS?

- A. The ability to scale computing resources up easily, with minimal friction and down with latency.
- B. The ability to scale computing resources up and down easily, with minimal friction.
- C. The ability to provision cloud computing resources in expectation of future demand.
- D. The ability to recover from business continuity events with minimal friction

**Answer: B**

#### NEW QUESTION 268

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 269

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 272

.....

## 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](#)