![PDF]

Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

NEW QUESTION 1
You want to use AWS CodeDeploy to deploy an application to Amazon EC2 instances running within an Amazon Virtual Private Cloud (VPC). What criterion must be met for this to be possible?

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

Answer: C

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

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

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

Answer: A

Explanation:
The maximum number of VPCs allowed per region is 5.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Appendix_Limits.html

NEW QUESTION 4
An organization is setting a website on the AWS VPC. The organization has blocked a few IPs to avoid a D-DOS attack. How can the organization configure that a request from the above mentioned IPs does not access the application instances?

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

Answer: D

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. AWS provides two features that the user can use to increase security in VPC: security groups and network ACLs. Security group works at the instance level while ACL works at the subnet level. ACL allows both allow and deny rules.
Thus, when the user wants to reject traffic from the selected IPs it is recommended to use ACL with subnets.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_ACLs.html

NEW QUESTION 5
In which step of using AWS Direct Connect should the user determine the required port speed?

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

Answer: D

Explanation:

To submit an AWS Direct Connect connection request, you need to provide the following information: Your contact information.
The AWS Direct Connect Location to connect to.
Details of AWS Direct Connect partner if you use the AWS Partner Network (APN) service. The port speed you require, either 1 Gbps or 10 Gbps.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/getstarted.html#ConnectionRequest

NEW QUESTION 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/IANI/latest/UserGuide/LimitationsOnEntities.html

NEW QUESTION 7
A user has created a NIySQL 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 8
In Amazon EIastiCache, 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 EIastiCache 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 capacit

Answer: B

Explanation:
In Amazon EIastiCache, 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 EIastiCache 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/AmazonEIastiCache/Iatest/UserGuide/CacheNode.Memcached.html

NEW QUESTION 9
A user is configuring MySQL RDS with PIOPS. What should be the minimum size of DB storage provided by the user?

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

Answer: D

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

NEW QUESTION 10
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-I7",
"Statement": [{
"Effect": "A||ow", "Action": ["iam:AddUserToGroup",
"iam:RemoveUserFromGroup", "iam:GetGroup"
]!
"Resource": "arn:aws:iam:: I23412341234:group/TestingGroup"

}I

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-I7",
"Statement": [{
"Effect": "A||ow", "Action": ["iam:AddUserToGroup",
"iam:RemoveUserFromGroup", "iam:GetGroup"
]!
"Resource": "arn:aws:iam:: 123412341234:group/ TestingGroup "
}I
Reference:
http://docs.aws.amazon.com/IAM/latest/UserGuide/CredentiaIs-Permissions-examples.htm|#creds-po|ici es-credentials

NEW QUESTION 10
One of your AWS Data Pipeline actMties 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/datapipe|ine/faqs/

NEW QUESTION 12
An organization is setting up a backup and restore system in AWS of their in premise system. The organization needs High AvaiIability(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 sewers 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 AvaiIability(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.cIoudfront.net/AWS_Disaster_Recovery.pdf

NEW QUESTION 15
Which of the following components of AWS Data Pipeline specifies the business logic of your data management?

A. Task Runner
B. Pipeline definition
C. AWS Direct Connect
D. Amazon Simple Storage Service (Amazon S3)

Answer: B

Explanation:
A pipeline definition specifies the business logic of your data management.
Reference: http://docs.aws.amazon.com/datapipeline/latest/DeveIoperGuide/what-is-datapipeline.htmI

NEW QUESTION 20
A user is planning to use EBS for his DB requirement. The user already has an EC2 instance running in the VPC private subnet. How can the user attach the EBS volume to a running instance?

A. The user can create EBS in the same zone as the subnet of instance and attach that EBS to instance.
B. It is not possible to attach an EBS to an instance running in VPC until the instance is stopped.
C. The user can specify the same subnet while creating EBS and then attach it to a running instance.
D. The user must create EBS within the same VPC and then attach it to a running instance.

Answer: A

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. 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. The instance launched will always be in the same availability zone of the respective subnet. When creating an EBS the user cannot specify the subnet or VPC. However, the user must create the EBS in the same zone as the instance so that it can attach the EBS volume to the running instance.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html#VPCSubnet

NEW QUESTION 23
An organization is purchasing licensed software. The software license can be registered only to a specific MAC Address. The organization is going to host the software in the AWS environment. How can the organization fulfil the license requirement as the MAC address changes every time an instance is started/stopped/terminated?

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

Answer: C

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. An Elastic Network Interface (ENI) is a virtual network interface that the user can attach to an instance in a VPC. An ENI can include attributes such as: a primary private IP address, one or more secondary private IP addresses, one elastic IP address per private IP address, one public IP address, one or more security groups, a MAC address, a source/destination check flag, and a description.
The user can create a network interface, attach it to an instance, detach it from an instance, and attach it to another instance. The attributes of a network interface follow the network interface as it is attached or detached from an instance and reattached to another instance. Thus, the user can maintain a fixed MAC using the network interface.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

NEW QUESTION 26
A user is trying to create a PIOPS EBS volume with 3 GB size and 90 IOPS. Will AWS create the volume?

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

Answer: C

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

NEW QUESTION 27
You have been given the task to define multiple AWS Data Pipeline schedules for different actMties 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 31
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 35
An EC2 instance that performs source/destination checks by default is launched in a private VPC subnet. All security, NACL, and routing definitions are configured as expected. A custom NAT instance is launched.
Which of the following must be done for the custom NAT instance to work?

A. The source/destination checks should be disabled on the NAT instance.
B. The NAT instance should be launched in public subnet.
C. The NAT instance should be configured with a public IP address.
D. The NAT instance should be configured with an elastic IP addres

Answer: A

Explanation:
Each EC2 instance performs source/destination checks by default. This means that the instance must be the source or destination of any traffic it sends or receives. However, a NAT instance must be able to send and receive traffic when the source or destination is not itself. Therefore, you must disable source/destination checks on the NAT instance.
Reference:
http://docs.aws.amazon.com/AmazonVPC/Iatest/UserGuide/VPC_NAT_|nstance.htm|#EIP_Disab|e_Src DestCheck

NEW QUESTION 39
An organization has created multiple components of a single application for compartmentalization. Currently all the components are hosted on a single EC2 instance. Due to security reasons the organization wants to implement two separate SSLs for the separate modules although it is already using VPC. How can the organization achieve this with a single instance?

A. You have to launch two instances each in a separate subnet and allow VPC peering for a single IP.
B. Create a VPC instance which will have multiple network interfaces with multiple elastic IP addresses.
C. Create a VPC instance which will have both the ACL and the security group attached to it and have separate rules for each IP address.
D. Create a VPC instance which will have multiple subnets attached to it and each will have a separate IP address.

Answer: B

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. With VPC the user can specify multiple private IP addresses for his instances.
The number of network interfaces and private IP addresses that a user can specify for an instance depends on the instance type. With each network interface the organization can assign an EIP. This scenario helps when the user wants to host multiple websites on a single EC2 instance by using multiple SSL certificates on a single server and associating each certificate with a specific EIP address. It also helps in scenarios for operating network appliances, such as firewalls or load balancers that have multiple private IP addresses for each network interface.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/NlultipleIP.htmI

NEW QUESTION 42
An organization is setting up a multi-site solution where the application runs on premise as well as on AWS to achieve the minimum recovery time objective(RTO). Which of the below mentioned configurations will not meet the requirements of the multi-site solution scenario?

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

Answer: C

Explanation:
AWS has many solutions for DR(Disaster recovery) and HA(High Availability). When the organization wants to have HA and DR with multi-site solution, it should setup two sites: one on premise and the other on AWS with full capacity. The organization should setup a weighted DNS service which can route traffic to both sites based on the weightage. When one of the sites fails it can route the entire load to another site. The organization would have minimal RTO in this scenario. If the organization setups a single DB instance, it will not work well in failover.
Instead they should have two separate DBs in each site and setup data replication based on RTO(recovery time objective)of the organization.
Reference: http://d36cz9buwru1tt.cloudfront.net/AWS_Disaster_Recovery.pdf

NEW QUESTION 43
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. MultipIe roles are assigned to an EC2 instance at a time.
D. Only one role can be assigned to an EC2 instance at a tim

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.htmI

NEW QUESTION 44
When using string conditions within IAM, short versions of the available comparators can be used instead of the more verbose ones.

streqi is the short version of the string condition.

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

Answer: A

Explanation:
When using string conditions within IANI, short versions of the available comparators can be used instead of the more verbose versions. For instance, streqi is the short version of StringEquals|gnoreCase 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 45
Select the correct statement about Amazon ElastiCache.

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

Answer: A

Explanation:
ElastiCache is a web service that makes it easy to set up, manage, and scale a distributed in-memory cache environment in the cloud. It provides a high-performance, scalable, and cost-effective caching solution, while removing the complexity associated with deploying and managing a distributed cache environment. With ElastiCache, you can quickly deploy your cache environment, without having to provision hardware or install software.
Reference: http://docs.aws.amazon.com/AmazonE|astiCache/latest/UserGuide/Whatls.html

NEW QUESTION 48
Which of the following cannot be done using AWS Data Pipeline?

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

Answer: C

Explanation:
AWS Data Pipeline is a web service that helps you reliably process and move data between different AWS compute and storage services as well as on-premise data sources at specified intervals. With AWS Data Pipeline, you can regularly access your data where it's stored, transform and process it at scale, and efficiently transfer the results to another AWS.
AWS Data Pipeline helps you easily create complex data processing workloads that are fault tolerant, repeatable, and highly available. AWS Data Pipeline also allows you to move and process data that was
previously locked up in on-premise data silos. Reference: http://aws.amazon.com/datapipe|ine/

NEW QUESTION 51
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 Nlanagement (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 resourc

Answer: C

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

NEW QUESTION 52
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 55
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 theusen
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 usen
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 59
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 61
How does AWS Data Pipeline execute actMties on on-premise resources or AWS resources that you manage?

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

Answer: A

Explanation:
To enable running actMties using on-premise resources, AWS Data Pipeline does the following: It supply a Task Runner package that can be installed on your on-premise hosts.
This package continuously polls the AWS Data Pipeline service for work to perform.
When it's time to run a particular actMty on your on-premise resources, it will issue the appropriate command to the Task Runner.
Reference: https://aws.amazon.com/datapipe|ine/faqs/

NEW QUESTION 66
In AWS IAM, which of the following predefined policy condition keys checks how long ago (in seconds) the MFA-validated security credentials making the request were issued using multi-factor authentication
(MFA)?

A. aws:MuItiFactorAuthAge
B. aws:Mu|tiFactorAuthLast
C. aws:MFAAge
D. aws:MuItiFactorAuthPrevious

Answer: A

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

NEW QUESTION 67
An organization is planning to host a web application in the AWS VPC. The organization does not want to host a database in the public cloud due to statutory requirements. How can the organization setup in this scenario?

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

Answer: A

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account.
The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to
connect with his data centre. When the user has configured this setup with Wizard, it will create a virtual private gateway to route all the traffic of the VPN subnet.
If the virtual private gateway is attached with VPC and the user deletes the VPC from the console it will first automatically detach the gateway and only then delete the VPC.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

NEW QUESTION 68
Do you need to use Amazon Cognito to use the Amazon Mobile Analytics service?

A. N
B. However, it is recommend by AWS to use Amazon Cognito for security best practices.
C. Ye
D. You need to use it only if you have IAM root access.
E. N
F. You cannot use it at all, and you need to use AWS IAM accounts.
G. Ye
H. It is recommended by AWS to use Amazon Cognito to use Amazon Mobile Analytics servic

Answer: A

Explanation:
You can initialize Amazon Mobile Analytics using AWS IAM accounts. AWS recommend using Amazon Cognito for security best practices.
Reference: http://aws.amazon.com/mobi|eanalytics/faqs/

NEW QUESTION 71
Which of the following AWS services can be used to define alarms to trigger on a certain actMty, such as actMty success, failure, or delay in AWS Data Pipeline?

A. Amazon SES
B. Amazon CodeDepIoy
C. Amazon SNS
D. Amazon SQS

Answer: C

Explanation:
In AWS Data Pipeline, you can define Amazon SNS alarms to trigger on actMties such as success, failure, or delay by creating an alarm object and referencing it in the onFail, onSuccess, or onLate slots of the actMty object.
Reference: https://aws.amazon.com/datapipe|ine/faqs/

NEW QUESTION 72
Which of the following components of AWS Data Pipeline polls for tasks and then performs those tasks?

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

Answer: B

Explanation:
Task Runner polls for tasks and then performs those tasks.
Reference: http://docs.aws.amazon.com/datapipeline/latest/DeveIoperGuide/what-is-datapipeline.html

NEW QUESTION 76
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/AmazonEIastiCache/latest/UserGuide/GettingStarted.AuthorizeAccess.htm|

NEW QUESTION 78
A read only news reporting site with a combined web and application tier and a database tier that receives large and unpredictable traffic demands must be able to respond to these traffic fluctuations automatically. What AWS services should be used meet these requirements?

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

Answer: A

NEW QUESTION 81
You are tasked with moving a legacy application from a virtual machine running Inside your datacenter to an Amazon VPC Unfortunately this app requires access to a number of on-premises services and no one who configured the app still works for your company. Even worse there's no documentation for it.
What will allow the application running inside the VPC to reach back and access its internal dependencies
without being reconfigured? (Choose 3 answers)

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

Answer: ADF

NEW QUESTION 86
You have a periodic Image analysis application that gets some files In Input analyzes them and tor each file writes some data in output to a ten file the number of files in input per day is high and concentrated in a few hours of the day.
Currently you have a server on EC2 with a large EBS volume that hosts the input data and the results it takes almost 20 hours per day to complete the process
What services could be used to reduce the elaboration time and improve the availability of the solution?

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

Answer: D

NEW QUESTION 89
Your department creates regular analytics reports from your company's log files All log data is collected in Amazon S3 and processed by daily Amazon Elastic MapReduce (EMR) jobs that generate daily PDF reports and aggregated tables in CSV format for an Amazon Redshift data warehouse.
Your CFO requests that you optimize the cost structure for this system.
Which of the following alternatives will lower costs without compromising average performance of the system or data integrity for the raw data?

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

Answer: C

NEW QUESTION 94
Your company runs a customer facing event registration site This site is built with a 3-tier architecture with web and application tier servers and a MySQL database The application requires 6 web tier sewers and 6 application tier servers for normal operation, but can run on a minimum of 65% server capacity and a single NIySQL database. When deploying this application in a region with three availability zones (AZs) which architecture provides high availability?

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

D. A web tier deployed across 3 AZs with 2 EC2 (Elastic Compute Cloud) instances in each AZ Inside an Auto Scaling Group behind an ELB (elastic load balancer). And an application tier deployed across 3 AZs with 2 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and a Multi-AZ RDS (Relational Database services) deployment.

Answer: D

NEW QUESTION 95
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 applicatio
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 zone
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 applicatio
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 loa
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 97
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 VNIDK is almost full; Nle virtual network interface still uses the 10IV|bps 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 me 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 me ec2-bundle-instance API to Import an Image of the VNI into EC2

Answer: A

NEW QUESTION 102
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 DynamoDS 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 105
Your company is getting ready to do a major public announcement of a social media site on AWS. The website is running on EC2 instances deployed across multiple Availability Zones with a Multi-AZ RDS MySQL Extra Large DB Instance. The site performs a high number of small reads and writes per second and relies on an eventual consistency model. After comprehensive tests you discover that there is read contention on RDS MySQL. Which are the best approaches to meet these requirements? (Choose 2 answers)

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

Answer: AC

NEW QUESTION 108
Refer to the architecture diagram above of a batch processing solution using Simple Queue Service (SQS) to set up a message queue between EC2 instances which are used as batch processors Cloud Watch monitors the number of Job requests (queued messages) and an Auto Scaling group adds or deletes batch sewers automatically based on parameters set in Cloud Watch alarms. You can use this architecture to implement which of the following features in a cost effective and efficient manner?

A. Reduce the overall lime for executing jobs through parallel processing by allowing a busy EC2 instance that receives a message to pass it to the next instance

in a daisy-chain setup.
B. Implement fault tolerance against EC2 instance failure since messages would remain in SQS and worn can continue with recovery of EC2 instances implement fault tolerance against SQS failure by backing up messages to S3.
C. Implement message passing between EC2 instances within a batch by exchanging messages throughSQS.
D. Coordinate number of EC2 instances with number of job requests automatically thus Improving cost effectiveness.
E. Handle high priority jobs before lower priority jobs by assigning a priority metadata field to SQS messages.

Answer: D

NEW QUESTION 112
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 me web application rapidly using CloudFormation.
The objective is to minimize changes to the existing web application, control the throughput of DynamoDB used for the synchronization of data and synchronize only the modified elements.
Which design would you choose to meet these requirements?

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

Answer: A

NEW QUESTION 115
You are designing a social media site and are considering how to mitigate distributed denial-of-service (DDoS) attacks. Which of the below are viable mitigation techniques? (Choose 3 answers)

A. Add multiple elastic network interfaces (ENIs) to each EC2 instance to increase the network bandwidth.
B. Use dedicated instances to ensure that each instance has the maximum performance possible.
C. Use an Amazon C|oudFront distribution for both static and dynamic content.
D. Use an Elastic Load Balancer with auto scaling groups at the we
E. App and Amazon Relational Database Service (RDS) tiers
F. Add alert Amazon CloudWatch to look for high Network in and CPU utilization.
G. Create processes and capabilities to quickly add and remove rules to the instance OS firewall

Answer: CEF

NEW QUESTION 118
You must architect the migration of a web application to AWS. The application consists of Linux web servers running a custom web server. You are required to save the logs generated from the application to a durable location.
What options could you select to migrate the application to AWS? (Choose 2)

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

Answer: AD

NEW QUESTION 120
A customer has established an AWS Direct Connect connection to AWS. The link is up and routes are being advertised from the customer's end, however the customer is unable to connect from EC2 instances inside its VPC to servers residing in its datacenter.
Which of the following options provide a viable solution to remedy this situation? (Choose 2 answers)

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

Answer: AC

NEW QUESTION 122

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

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

Answer: BC

NEW QUESTION 124

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

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

Answer: B

NEW QUESTION 128

You are designing an SSUTLS solution that requires HTTPS clients to be authenticated by the Webserver using client certificate authentication. The solution must be resilient.
Which of the following options would you consider for configuring the web server infrastructure? (Choose 2 answers)

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

Answer: AB

NEW QUESTION 129

You are migrating a legacy client-server application to AWS. The application responds to a specific DNS domain (e.g. www.examp|e.com) and has a 2-tier architecture, with multiple application sewers and a database sewer. Remote clients use TCP to connect to the application servers. The application servers need to know the IP address of the clients in order to function properly and are currently taking that information from the TCP socket. A Multi-AZ RDS MySQL instance will be used for the database. During the migration you can change the application code, but you have to file a change request.
How would you implement the architecture on AWS in order to maximize scalability and high availability?

A. File a change request to implement Alias Resource support in the applicatio
B. Use Route 53 Alias Resource Record to distribute load on two application servers in different Azs.
C. File a change request to implement Latency Based Routing support in the applicatio
D. Use Route 53 with Latency Based Routing enabled to distribute load on two application servers in different Azs.
E. File a change request to implement Cross-Zone support in the applicatio
F. Use an ELB with a TCP Listener and Cross-Zone Load Balancing enabled, two application servers in different AZs.
G. File a change request to implement Proxy Protocol support in the applicatio
H. Use an ELB with a TCP Listener and Proxy Protocol enabled to distribute load on two application servers in different Azs.

Answer: D

NEW QUESTION 131

You are designing a personal document-archMng solution for your global enterprise with thousands of employee. Each employee has potentially gigabytes of data to be backed up in this archMng solution. The solution will be exposed to the employees as an application, where they can just drag and drop their files to the archMng system. Employees can retrieve their archives through a web interface. The corporate network has high bandwidth AWS Direct Connect connectMty to AWS.
You have a regulatory requirement that all data needs to be encrypted before being uploaded to the cloud.
How do you implement this in a highly available and cost-efficient way?

A. Manage encryption keys on-premises in an encrypted relational databas
B. Set up an on-premises server with sufficient storage to temporarily store files, and then upload them to Amazon S3, providing a client-side master key.
C. Mange encryption keys in a Hardware Security Module (HSM) appliance on-premises serve r with sufficient storage to temporarily store, encrypt, and upload files directly into Amazon Glacier.
D. Nlanage encryption keys in Amazon Key Management Service (KMS), upload to Amazon Simple Storage Service (S3) with client-side encryption using a KMS customer master key ID, and configure Amazon S3 lifecycle policies to store each object using the Amazon Glacier storage tier.

E. Manage encryption keys in an AWS C|oudHSNI applianc
F. Encrypt files prior to uploading on the employee desktop, and then upload directly into Amazon Glacier.

Answer: C

NEW QUESTION 135
You are responsible for a web application that consists of an Elastic Load Balancing (ELB) load balancer in front of an Auto Scaling group of Amazon Elastic Compute Cloud (EC2) instances. For a recent deployment of a new version of the application, a new Amazon Machine Image (AMI) was created, and the Auto Scaling group was updated with a new launch configuration that refers to this new AMI. During the deployment, you received complaints from users that the website was responding with errors. All instances passed the ELB health checks.
What should you do in order to avoid errors for future deployments? (Choose 2 answer)

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

Answer: CD

NEW QUESTION 139
You have an application running on an EC2 instance which will allow users to download files from a private S3 bucket using a pre-signed URL. Before generating the URL, the application should verify the existence of the file in S3. How should the application use AWS credentials to access the S3 bucket securely?

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

Answer: B

NEW QUESTION 142
A 3-tier e-commerce web application is current deployed on-premises and will be migrated to AWS for greater scalability and elasticity The web server currently shares read-only data using a network distributed file system The app server tier uses a clustering mechanism for discovery and shared session state that depends on IP multicast The database tier uses shared-storage clustering to provide database fall over capability, and uses several read slaves for scaling Data on all servers and the distributed file system directory is backed up weekly to off-site tapes
Which AWS storage and database architecture meets the requirements of the application?

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

Answer: C

NEW QUESTION 147
You need a persistent and durable storage to trace call actMty of an IVR (Interactive Voice Response) system. Call duration is mostly in the 2-3 minutes timeframe. Each traced call can be either active or terminated. An external application needs to know each minute the list of currently active calls. Usually there are a few calls/second, but once per month there is a periodic peak up to 1000 calls/second for a few hours. The system is open 24/7 and any downtime should be avoided. Historical data is periodically archived to files. Cost saving is a priority for this project.
What database implementation would better fit this scenario, keeping costs as low as possible?

A. Use DynamoDB with a "Calls" table and a Global Secondary Index on a "State" attribute that can equal to "active" or "terminated". In this way the Global Secondary Index can be used for all items in the table.
B. Use RDS Multi-AZ with a "CALLS" table and an indexed "STATE" field that can be equal to "ACT|VE"or 'TERMINATED". In this way the SQL query is optimized by the use of the Index.
C. Use RDS Nlulti-AZ with two tables, one for "ACT|VE_CALLS" and one for "TERMINATED_CALLS". In this way the "ACTIVE_CALLS" table is always small and effective to access.
D. Use DynamoDB with a "Calls" table and a Global Secondary Index on a "Is Active" attribute that is present for active calls onl
E. In this way the Global Secondary Index is sparse and more effective.

Answer: C

NEW QUESTION 150
An administrator is using Amazon CloudFormation to deploy a three tier web application that consists of a web tier and application tier that will utilize Amazon DynamoDB for storage when creating the CloudFormation template which of the following would allow the application instance access to the DynamoDB tables without exposing API credentials?

A. Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and associate the Role to the application instances by referencing an instance profile.
B. Use the Parameter section in the Cloud Formation template to nave the user input Access and Secret Keys from an already created IAM user that has me permissions required to read and write from the required DynamoDB table.
C. Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and reference the Role in the instance profile property of the application instance.
D. Create an identity and Access Management user in the CloudFormation template that has permissions to read and write from the required DynamoDB table, use the GetAtt function to retrieve the Access and secret keys and pass them to the application instance through user-data.

Answer: C

NEW QUESTION 152
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-compliam identity provider (IDP) to retrieve temporary security credentials to enable NOC members to sign in to the AWS Management Console.

Answer: D

NEW QUESTION 154
You are developing a new mobile application and are considering storing user preferences in AWS.2w 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 dat
B. Deploy apublic 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 tabl
E. Utilize ST
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 prMlege 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' S3 objec
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 159
A web-startup runs its very successful social news application on Amazon EC2 with an Elastic Load Balancer, an Auto-Scaling group of Java/Tomcat application-servers, and DynamoDB as data store. The main web-application best runs on m2 x large instances since it is highly memory- bound Each new deployment requires semi-automated creation and testing of a new AMI for the application servers which takes quite a while ana is therefore only done once per week.
Recently, a new chat feature has been implemented in nodejs and wails to be integrated in the architecture. First tests show that the new component is CPU bound Because the company has some experience with using Chef, they decided to streamline the deployment process and use AWS Ops Works as an application life cycle tool to simplify management of the application and reduce the deployment cycles.
What configuration in AWS Ops Works is necessary to integrate the new chat module in the most cost-efficient and filexible way?

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

Answer: C

NEW QUESTION 162
Which of the following are characteristics of Amazon VPC subnets? Choose 2 answers

A. Each subnet spans at least 2 Availability Zones to provide a high-availability environment.
B. Each subnet maps to a single Availability Zone.
C. CIDR block mask of /25 is the smallest range supported.
D. By default, all subnets can route between each other, whether they are private or public.
E. Instances in a private subnet can communicate with the Internet only if they have an Elastic I

Answer: AE

NEW QUESTION 165
In AWS, which security aspects are the customer's responsibility? Choose 4 answers

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

Answer: ACDF

NEW QUESTION 167
You have recently joined a startup company building sensors to measure street noise and air quality in urban areas. The company has been running a pilot deployment of around 100 sensors for 3 months each sensor uploads 1KB of sensor data every minute to a backend hosted on AWS.
During the pilot, you measured a peak or 10 IOPS on the database, and you stored an =average of 3GB of sensor data per month in the database.
The current deployment consists of a load-balanced auto scaled Ingestion layer using EC2 instances and a PostgreSQL RDS database with 500GB standard storage.
The pilot is considered a success and your CEO has managed to get the attention or some potential investors. The business plan requires a deployment of at least 100K sensors which needs to be supported by the backend. You also need to store sensor data for at least two years to be able to compare year over year Improvements.
To secure funding, you have to make sure that the platform meets these requirements and leaves room for further scaling. Which setup win meet the requirements?

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

Answer: C

NEW QUESTION 170
Your team has a tomcat-based Java application you need to deploy into development, test and production environments. After some research, you opt to use Elastic Beanstalk due to its tight integration with your developer tools and RDS due to its ease of management. Your QA team lead points out that you need to roll a sanitized set of production data into your environment on a nightly basis.
Similarly, other software teams in your org want access to that same restored data via their EC2 instances in your VPC .The optimal setup for persistence and security that meets the above requirements would be the following.

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

Answer: A

NEW QUESTION 173
Your company has an on-premises multi-tier PHP web application, which recently experienced downtime due to a large burst In web traffic due to a company announcement Over the coming days, you are expecting similar announcements to drive similar unpredictable bursts, and are looking to find ways to quickly improve your infrastructures ability to handle unexpected increases in traffic.
The application currently consists of 2 tiers a web tier which consists of a load balancer and several Linux Apache web servers as well as a database tier which hosts a Linux server hosting a MySQL database. Which scenario below will provide full site functionality, while helping to improve the ability of your application in the short timeframe required?

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

Answer: C

NEW QUESTION 178
You are implementing AWS Direct Connect. You intend to use AWS public service end points such as Amazon S3, across the AWS Direct Connect link. You want other Internet traffic to use your existing link to an Internet Service Provider.
What is the correct way to configure AWS Direct connect for access to services such as Amazon S3?

A. Configure a public Interface on your AWS Direct Connect link Configure a static route via your AWS Direct Connect link that points to Amazon S3 Advertise a default route to AWS using BGP.

B. Create a private interface on your AWS Direct Connect lin
C. Configure a static route via your AWS Direct connect link that points to Amazon S3 Configure specific routes to your network in your VPC.
D. Create a public interface on your AWS Direct Connect link Redistribute BGP routes into your existing routing infrastructure; advertise specific routes for your network to AWS.
E. Create a private interface on your AWS Direct connect lin
F. Redistribute BGP routes into your existing routing infrastructure and advertise a default route to AWS.

Answer: C

NEW QUESTION 183
Your application is using an ELB in front of an Auto Scaling group of web/application sewers deployed across two AZs and a Multi-AZ RDS Instance for data persistence.
The database CPU is often above 80% usage and 90% of I/O operations on the database are reads. To improve performance you recently added a single-node Memcached ElastiCache Cluster to cache frequent DB query results. In the next weeks the overall workload is expected to grow by 30%.
Do you need to change anything in the architecture to maintain the high availability or the application with the anticipated additional load? Why?

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

Answer: A

NEW QUESTION 188
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 193
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": "A||ow",
"Resource": ["arn:aws:s3::zbucket-name"], "Condition":{"StringLike":{"s3:prefix":["home/${aws:username}/*"]}}
}!
{
"Action":["s3:*"], "Effect":"A||ow",
"Resource": ["arn:aws:s3:::bucket-name/home/${aws:username}/*"]
}
}

A. True
B. False

Answer: B

NEW QUESTION 195
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 199
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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](https://www.certshared.com/exam/AWS-Certified-Solutions-Architect-Professional/)