Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

https://www.2passeasy.com/dumps/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 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
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 gatewa

Answer: B

Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. The organization wants to extend their network into the cloud and also directly access the internet from their AWS VPC. Thus, the organization should setup a Virtual Private Cloud (VPC) with a public subnet and a private subnet, and a virtual private gateway to enable communication with their data center network over an IPsec VPN tunnel. To setup this configuration the organization needs to use the Amazon VPC with a VPN connection. The organization network administrator must designate a physical appliance as a customer gateway and configure it. The organization would need the below mentioned information to setup this configuration:
The type of customer gateway, such as Cisco ASA, Juniper J-Series, Juniper SSG, Yamaha Internet-routable IP address (static) of the customer gateway's external interface
Border Gateway Protocol (BGP) Autonomous System Number (ASN) of the customer gateway, if the organization is creating a dynamically routed VPN connection.
Internal network IP ranges that the user wants to advertise over the VPN connection to the VPC. Reference:
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html

NEW QUESTION 4
An organization is planning to host a Wordpress blog as well a joomla CMS on a single instance launched with VPC. The organization wants to have separate domains for each application and assign them using Route 53. The organization may have about ten instances each with two applications as mentioned above. While launching the instance, the organization configured two separate network interfaces (primary + ENI) and wanted to have two elastic IPs for that instance. It was suggested to use a public IP from AWS instead of an elastic IP as the number of elastic IPs is restricted. What action will you recommend to the organization?

A. I agree with the suggestion but will prefer that the organization should use separate subnets with each ENI for different public IPs.
B. I do not agree as it is required to have only an elastic IP since an instance has more than one ENI and AWS does not assign a public IP to an instance with multiple ENIs.
C. I do not agree as AWS VPC does not attach a public IP to an ENI; so the user has to use only an elastic IP only.
D. I agree with the suggestion and it is recommended to use a public IP from AWS since the organization is going to use DNS with Route 53.

Answer: B

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. An Elastic Network Interface (ENI) is a virtual network interface that the user can attach to an instance in a VPC.

The user can attach up to two ENIs with a single instance. However, AWS cannot assign a public IP when there are two ENIs attached to a single instance. It is recommended to assign an elastic IP in this scenario. If the organization wants more than 5 E|Ps they can request AWS to increase the number.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.htmI

NEW QUESTION 5
A customer has a website which shows all the deals available across the market. The site experiences a load of 5 large EC2 instances generally. However, a week before ThanksgMng vacation they encounter a load of almost 20 large instances. The load during that period varies over the day based on the office timings. Which of the below mentioned solutions is cost effective as well as help the website achieve better performance?

A. Setup to run 10 instances during the pre-vacation period and only scale up during the office time by launching 10 more instances using the AutoScaling schedule.
B. Keep only 10 instances running and manually launch 10 instances every day during office hours.
C. During the pre-vacation period setup 20 instances to run continuously.
D. During the pre-vacation period setup a scenario where the organization has 15 instances running and 5 instances to scale up and down using Auto Scaling based on the network I/O policy.

Answer: D

Explanation:
AWS provides an on demand, scalable infrastructure. AWS EC2 allows the user to launch On-Demand instances and the organization should create an AMI of the running instance. When the organization is experiencing varying loads and the time of the load is not known but it is higher than the routine traffic it is recommended that the organization launches a few instances before hand and then setups AutoScaling with policies which scale up and down as per the EC2 metrics, such as Network I/O or CPU utilization.
If the organization keeps all 10 additional instances as a part of the AutoScaling policy sometimes during a sudden higher load it may take time to launch instances and may not give an optimal performance. This is the reason it is recommended that the organization keeps an additional 5 instances running and the next 5 instances scheduled as per the AutoScaling policy for cost effectiveness.
Reference: http://media.amazonwebservices.com/AWS_Web_Hosting_Best_Practices.pdf

NEW QUESTION 6
An organization has 4 people in the IT operations team who are responsible to manage the AWS infrastructure. The organization wants to setup that each user will have access to launch and manage an instance in a zone which the other user cannot modify. Which of the below mentioned options is the best solution to set this up?

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

Answer: D

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. The VPC also work with IAM and the organization can create IAM users who have access to various VPC services. The organization can setup access for the IAM user who can modify the security groups of the VPC. The sample policy is given below:
{
"Version": "2012-10-I7",
"Statement":
[{ "Effect": "AIIow", "Action": "ec2:RunInstances", "Resource":
["arn:aws:ec2:region::image/ami-*", "arn:aws:ec2:region:accountzsubnet/subnet-1a2b3c4d", "arn:aws:ec2:region:account:network-interface/*",
"arn:aws:ec2:region:account:vo|ume/*", "arn:aws:ec2:region:account:key-pair/*", "arn:aws:ec2:region:account:security-group/sg-123abc123"]
}I I
With this policy the user can create four subnets in separate zones and provide IAM user access to each subnet
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_IANI.htmI

NEW QUESTION 7
You have subscribed to the AWS Business and Enterprise support plan. Your business has a backlog of problems, and you need about 20 of your IAM users to open technical support cases. How many users can open technical support cases under the AWS Business and Enterprise support plan?

A. 5 users
B. 10 users
C. Unlimited
D. 1 user

Answer: C

Explanation:
In the context of AWS support, the Business and Enterprise support plans allow an unlimited number of users to open technical support cases (supported by AWS Identity and Access Management (IAM)). Reference: https://aws.amazon.com/premiumsupport/faqs/

NEW QUESTION 8
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-f|ow/

NEW QUESTION 9
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 rang

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 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
How can multiple compute resources be used on the same pipeline in AWS Data Pipeline?

A. You can use multiple compute resources on the same pipeline by defining multiple cluster objects in your definition file and associating the cluster to use for each actMty via its runsOn field.
B. You can use multiple compute resources on the same pipeline by defining multiple cluster definition files.
C. You can use multiple compute resources on the same pipeline by defining multiple clusters for your actMty.
D. You cannot use multiple compute resources on the same pipelin

Answer: A

Explanation:
Multiple compute resources can be used on the same pipeline in AWS Data Pipeline by defining multiple cluster objects in your definition file and associating the cluster to use for each actMty via its runsOn field, which allows pipelines to combine AWS and on-premise resources, or to use a mix of instance types for their actMties.
Reference: https://aws.amazon.com/datapipe|ine/faqs/

NEW QUESTION 14
The MySecureData company has five branches across the globe. They want to expand their data centers such that their web server will be in the AWS and each branch would have their own database in the local data center. Based on the user login, the company wants to connect to the data center. How can MySecureData company implement this scenario with the AWS VPC?

A. Create five VPCs with the public subnet for the app server and setup the VPN gateway for each VPN to connect them indMdually.
B. Use the AWS VPN CIoudHub to communicate with multiple VPN connections.
C. Use the AWS CIoudGateway to communicate with multiple VPN connections.

D. It is not possible to connect different data centers from a single VPC.

Answer: B

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. If the organization has multiple VPN connections, he can provide secure communication between sites using the AWS VPN CloudHub.
The VPN CloudHub operates on a simple hub-and-spoke model that the user can use with or without a VPC. This design is suitable for customers with multiple branch offices and existing internet connections who would like to implement a convenient, potentially low-cost hub-and-spoke model for primary or backup connectMty between remote offices.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPN_CIoudHub.htmI

NEW QUESTION 16
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 18
True or False: In Amazon EIastiCache 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 EIastiCache, 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/AmazonEIastiCache/Iatest/UserGuide/Replication.Redis.Groups.htmI

NEW QUESTION 23
True or False: Amazon EIastiCache supports the Redis key-value store.

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

Answer: C

Explanation:
This is true. EIastiCache supports two open-source in-memory caching engines: 1. Memcached - a widely adopted memory object caching system. EIastiCache 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. EIastiCache supports Master / Slave replication and Multi-AZ which can be used to achieve cross AZ redundancy.
Reference: https://aws.amazon.com/eIasticache/

NEW QUESTION 28
Which of the following is NOT an advantage of using AWS Direct Connect?

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

Answer: A

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

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

NEW QUESTION 31
An organization is setting up an application on AWS to have both High Availabilty (HA) and Disaster Recovery (DR). The organization wants to have both Recovery point objective (RPO) and Recovery time objective (RTO) of 10 minutes. Which of the below mentioned service configurations does not help the organization achieve the said RPO and RTO?

A. Take a snapshot of the data every 10 minutes and copy it to the other region.
B. Use an elastic IP to assign to a running instance and use Route 53 to map the user's domain with that IP.
C. Create ELB with multi- region routing to allow automated failover when required.
D. Use an AMI copy to keep the AMI available in other region

Answer: C

Explanation:
AWS provides an on demand, scalable infrastructure. AWS EC2 allows the user to launch On-Demand instances and the organization should create an AMI of the running instance. Copy the AMI to another region to enable Disaster Recovery (DR) in case of region failure. The organization should also use EBS for persistent storage and take a snapshot every 10 minutes to meet Recovery time objective (RTO). They should also setup an elastic IP and use it with Route 53 to route requests to the same IP.
When one of the instances fails the organization can launch new instances and assign the same EIP to a new instance to achieve High Availability (HA). The ELB works only for a particular region and does not route requests across regions.
Reference: http://d36cz9buwru1tt.c|oudfront.net/AWS_Disaster_Recovery.pdf

NEW QUESTION 33
By default, what is the maximum number of Cache Nodes you can run in Amazon ElastiCache?

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

Answer: A

Explanation:
In Amazon ElastiCache, you can run a maximum of 20 Cache Nodes. Reference: http://aws.amazon.com/e|asticache/faqs/

NEW QUESTION 34
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.html

NEW QUESTION 38
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_credentiaIs_temp_control-access_disable-perms.h tml

NEW QUESTION 43
An organization is planning to setup a management network on the AWS VPC. The organization is trying to secure the webserver on a single VPC instance such that it allows the internet traffic as well as the back-end management traffic. The organization wants to make so that the back end management network interface can receive the SSH traffic only from a selected IP range, while the internet facing webserver will have an IP address which can receive traffic from all the internet IPs.
How can the organization achieve this by running web server on a single instance?

A. It is not possible to have two IP addresses for a single instance.

B. The organization should create two network interfaces with the same subnet and security group to assign separate IPs to each network interface.
C. The organization should create two network interfaces with separate subnets so one instance can have two subnets and the respective security groups for controlled access.
D. The organization should launch an instance with two separate subnets using the same network interface which allows to have a separate CIDR as well as security groups.

Answer: C

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. An Elastic Network Interface (ENI) is a virtual network interface that the user can attach to an instance in a VPC.
The user can create a management network using two separate network interfaces. For the present scenario it is required that the secondary network interface on the instance handles the public facing traffic and the primary network interface handles the back-end management traffic and it is connected to a separate subnet in the VPC that has more restrictive access controls. The public facing interface, which may or may not be behind a load balancer, has an associated security group to allow access to the server from the internet while the private facing interface has an associated security group allowing SSH access only from an allowed range of IP addresses either within the VPC or from the internet, a private subnet within the VPC or a virtual private gateway.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

NEW QUESTION 46
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 49
Cognito Sync is an AWS service that you can use to synchronize user profile data across mobile devices without requiring your own backend. When the device is online, you can synchronize data. If you also set up push sync, what does it allow you to do?

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

Answer: C

Explanation:
Cognito Sync is an AWS service that you can use to synchronize user profile data across mobile devices without requiring your own backend. When the device is online, you can synchronize data, and if you have
also set up push sync, notify other devices immediately that an update is available. Reference: http://docs.aws.amazon.com/cognito/devguide/sync/

NEW QUESTION 54
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 onl

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 58
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 61
A bucket owner has allowed another account's IAM users to upload or access objects in his bucket. The IAM user of Account A is trying to access an object created by the IAM user of account B. What will happen in this scenario?

A. It is not possible to give permission to multiple IAM users
B. AWS S3 will verify proper rights given by the owner of Account A, the bucket owner as well as by the IAM user B to the object
C. The bucket policy may not be created as S3 will give error due to conflict of Access Rights
D. It is not possible that the IAM user of one account accesses objects of the other IAM user

Answer: B

Explanation:
If a IAM user is trying to perform some action on an 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.htmI

NEW QUESTION 63
An organization has hosted an application on the EC2 instances. There will be multiple users connecting to the instance for setup and configuration of application. The organization is planning to implement certain security best practices. Which of the below mentioned pointers will not help the organization achieve better security arrangement?

A. Allow only IAM users to connect with the EC2 instances with their own secret access key.
B. Create a procedure to revoke the access rights of the indMdual user when they are not required to connect to EC2 instance anymore for the purpose of application configuration.
C. Apply the latest patch of OS and always keep it updated.
D. Disable the password based login for all the user
E. All the users should use their own keys to connect with the instance securely.

Answer: A

Explanation:
Since AWS is a public cloud any application hosted on EC2 is prone to hacker attacks. It becomes extremely important for a user to setup a proper security mechanism on the EC2 instances. A few of the security measures are listed below:
Always keep the OS updated with the latest patch
Always create separate users with in OS if they need to connect with the EC2 instances, create their keys and disable their password
Create a procedure using which the admin can revoke the access of the user when the business work on the EC2 instance is completed
Lock down unnecessary ports
Audit any proprietary applications that the user may be running on the EC2 instance
Provide temporary escalated prMleges, such as sudo for users who need to perform occasional prMleged tasks
The IAM is useful when users are required to work with AWS resources and actions, such as launching an instance. It is not useful to connect (RDP / SSH) with an instance.
Reference: http://aws.amazon.com/articles/1233/

NEW QUESTION 67
In Amazon SNS, to send push notifications to mobile devices using Amazon SNS and ADM, you need to obtain the following, except:

A. Device token
B. Client ID
C. Registration ID
D. Client secret

Answer: A

Explanation:
To send push notifications to mobile devices using Amazon SNS and ADM, you need to obtain the following: Registration ID and Client secret.
Reference: http://docs.aws.amazon.com/sns/latest/dg/SNSMobi|ePushPrereq.htmI

NEW QUESTION 71
True or False : "|n the context of Amazon ElastiCache, from the application's point of view, connecting to the cluster configuration endpoint is no different than connecting directly to an indMdual cache node."

A. True, from the application's point of view, connecting to the cluster configuration endpoint is no different than connecting directly to an indMdual cache node since, each has a unique node identifier.
B. True, from the application's point of view, connecting to the cluster configuration endpoint is no different than connecting directly to an indMdual cache node.
C. False, you can connect to a cache node, but not to a cluster configuration endpoint.
D. False, you can connect to a cluster configuration endpoint, but not to a cache nod

Answer: B

Explanation:

This is true. From the application's point of view, connecting to the cluster configuration endpoint is no different than connecting directly to an indMdual cache node. In the process of connecting to cache nodes, the application resolves the configuration endpoint's DNS name. Because the configuration endpoint maintains CNAME entries for all of the cache nodes, the DNS name resolves to one of the nodes; the client can then connect to that node.
Reference: http://docs.aws.amazon.com/AmazonEIastiCache/latest/UserGuide/AutoDiscovery.HowAutoDiscoveryW orks.htmI

NEW QUESTION 76
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 78
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 80
Identify an application that polls AWS Data Pipeline for tasks and then performs those tasks.

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

Answer: C

Explanation:
A task runner is an application that polls AWS Data Pipeline for tasks and then performs those tasks. You can either use Task Runner as provided by AWS Data Pipeline, or create a custom Task Runner application.
Task Runner is a default implementation of a task runner that is provided by AWS Data Pipeline. When Task Runner is installed and configured, it polls AWS Data Pipeline for tasks associated with pipelines that you have activated. When a task is assigned to Task Runner, it performs that task and reports its status back to AWS Data Pipeline. If your workflow requires non-default behavior, you'II need to implement that functionality in a custom task runner.
Reference:
http://docs.aws.amazon.com/datapipeline/latest/DeveIoperGuide/dp-how-remote-taskrunner-client.html

NEW QUESTION 81
Regarding Identity and Access Management (IAM), Which type of special account belonging to your application allows your code to access Google services programmatically?

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

Answer: A

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

Reference: https://cloud.google.com/iam/docs/service-accounts

NEW QUESTION 86
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 88
ExamKil|er 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. ExamKil|er wants to setup a centralized server which can connect with the LMS of each tenant upgrade if required. ExamKil|er also wants to ensure that one tenant VPC should not be able to connect to the other tenant VPC for security reasons. How can ExamKil|er setup this scenario?

A. ExamKil|er has to setup one centralized VPC which will peer in to all the other VPCs of the tenants.
B. ExamKil|er 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. ExamKil|er 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. ExamKil|er 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/Iatest/PeeringGuide/peering-configurations-full-access.html# many-vpcs-full-acces

NEW QUESTION 91
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 96
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 rol

Answer: A

Explanation:
Amazon designed IANI 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 100
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 103
In AWS IAM, which of the following predefined policy condition keys checks how long ago (in seconds) the MFA-validated security credentials making the request were issued using multi-factor authentication
(MFA)?

A. aws: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.htmI

NEW QUESTION 104
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/Iatest/UserGuide/VPC_Subnets.html

NEW QUESTION 105
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 subne

Answer: A

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. The VPC is always specific to a region. The user can create a VPC which can span multiple Availability Zones by adding one or more subnets in each Availability Zone. Each subnet must reside entirely within one Availability Zone and cannot span across zones.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html#VPCSubnet

NEW QUESTION 110
A user is creating a PIOPS volume. What is the maximum ratio the user should configure between PIOPS and the volume size?

A. 5
B. 10
C. 20
D. 30

Answer: D

Explanation:

Provisioned IOPS volumes are designed to meet the needs of I/O-intensive workloads, particularly database workloads that are sensitive to storage performance and consistency in random access I/O throughput. A provisioned IOPS volume can range in size from 10 GB to 1 TB and the user can provision up to 4000 IOPS per volume.

The ratio of IOPS provisioned to the volume size requested can be a maximum of 30; for example, a volume with 3000 IOPS must be atleast 100 GB.

Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVo|umeTypes.htm|

NEW QUESTION 112
What is a possible reason you would need to edit claims issued in a SAML token?

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

Answer: A

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

NEW QUESTION 115
You're trying to delete an SSL certificate from the IAM certificate store, and you're getting the message "Certificate: <certificate-id> is being used by CloudFront."
Which of the following statements is probably the reason why you are getting this error?

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

Answer: C

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

NEW QUESTION 116
A user has set the IAM policy where it denies all requests if a request is not from IP 10.10.10.1/32. The other policy says allow all requests between 5 PM to 7 PM.
What will happen when a user is requesting access from IP 55.109.10.12/32 at 6 PM?

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

Answer: A

Explanation:
When a request is made, the AWS IAM policy decides whether a given request should be allowed or denied. The evaluation logic follows these rules:
By default, all requests are denied. (In general, requests made using the account credentials for resources in the account are always allowed.)
An explicit allow policy overrides this default.
An explicit deny policy overrides any allows.
In this case since there are explicit deny and explicit allow statements. Thus, the request will be denied since deny overrides allow.
Reference: http://docs.aws.amazon.com/IAM/Iatest/UserGuide/AccessPoIicyLanguage_EvaIuationLogic.html

NEW QUESTION 117
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|eanaIytics/faqs/

NEW QUESTION 120
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 onFaiI, onSuccess, or onLate slots of the actMty object.
Reference: https://aws.amazon.com/datapipe|ine/faqs/

NEW QUESTION 121
Mike is appointed as Cloud Consultant in ExamKi|Ier.com. ExamKiI|er has the following VPCs set-up in the US East Region:
A VPC with CIDR block 10.10.0.0/16, a subnet in that VPC with CIDR block 10.10.1.0/24 A VPC with CIDR block 10.40.0.0/16, a subnet in that VPC with CIDR block 10.40.1.0/24
ExamKiIIer.com is trying to establish network connection between two subnets, a subnet with CIDR block 10.10.1.0/24 and another subnet with CIDR block 10.40.1.0/24. Which one of the following solutions should IV|ike recommend to ExamKiI|er.com?

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

Answer: C

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

NEW QUESTION 125
An organization, which has the AWS account ID as Q99988887777, has created 50 IAM users. All the users are added to the same group examkiller. If the organization has enabled that each IAM user can login with the AWS console, which AWS login URL will the IAM users use??

A. https://Q99988887777.aws.amazon.com/examkiIIer/
B. https://signin.aws.amazon.com/examki|Ier/
C. https://examkiller.signin.aws.amazon.com/999988887777/consoIe/
D. https://999988887777.signin.aws.amazon.com/consoIe/

Answer: D

Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. Once the organization has created the IAM users, they will have a separate AWS console URL to login to the AWS console. The console login URL for the IAM user will be https:// AWS_Account_ID.signin.aws.amazon.com/consoIe/. It uses only the AWS account ID and does not depend on the group or user ID.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccountAlias.htmI

NEW QUESTION 129
You're running an application on-premises due to its dependency on non-x86 hardware and want to use AWS for data backup. Your backup application is only able to write to POSIX-compatible block-based storage. You have 140TB of data and would like to mount it as a single folder on your file server Users must be able to access portions of this data while the backups are taking place. What backup solution would be most appropriate for this use case?

A. Use Storage Gateway and configure it to use Gateway Cached volumes.
B. Configure your backup software to use S3 as the target for your data backups.
C. Configure your backup software to use Glacier as the target for your data backups.
D. Use Storage Gateway and configure it to use Gateway Stored volume

Answer: A

NEW QUESTION 134
You have deployed a web application targeting a global audience across multiple AWS Regions under the domain name.exampIe.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. Dunning a DR test you notice that when you disable all web sewers 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 me disabled web sewers.
C. The value of the weight associated with the latency alias resource record set in the region with the disabled sewers is higher than the weight for the other region.
D. One of the two working web sewers 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 137
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 EIastiCache Memcached in an autoscaimg group monitored with CIoudWatch and RDS with read replicas.
B. Stateful instances for the web and application tier in an autoscaling group monitored with CIoudWatch and RDS with read replicas.
C. Stateful instances for the web and application tier in an autoscaling group monitored with CIoudWatc
D. And multi-AZ RDS.
E. Stateless instances for the web and application tier synchronized using EIastiCache Memcached in an autoscaling group monitored with CIoudWatch and multi-AZ RDS.

Answer: A

NEW QUESTION 142
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 145
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 indMdual 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 c3.8x|arge instance in us-east-1. Provision 4x1TB EBS volumes, attach them to the instance, and configure them as a single RAID 5 volum
B. Ensure that EBS snapshots are performed every 15 minutes.
C. Instantiate a c3.8xlarge instance in us-east-1. Provision 3xITB EBS volumes, attach them to the Instance, and configure them as a single RAID 0 volum
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 instanc
G. Provision 3x1TB EBS volumes, attach them to the instance, and configure them as a second RAID 0 volum
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 instanc
M. Configure synchronous, blocklevel replication to an identically configured instance in us-east-1b.

Answer: C

NEW QUESTION 147
You require the ability to analyze a large amount of data, which is stored on Amazon S3 using Amazon Elastic Map Reduce. You are using the cc2 8x large Instance type, whose CPUs are mostly idle during processing. Which of the below would be the most cost efficient way to reduce the runtime of the job?

A. Create more smaller flies on Amazon S3.
B. Add additional cc2 8x large instances by introducing a task group.
C. Use smaller instances that have higher aggregate I/O performance.
D. Create fewer, larger files on Amazon S3.

Answer: C

NEW QUESTION 151
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 clien
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 threshol
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 tor each customer with a Bucket Policy that permits access only to that one customer.

Answer: A

NEW QUESTION 152
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 154
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 156
A web company is looking to implement an external payment service into their highly available application deployed in a VPC Their application EC2 instances are behind a public lacing ELB Auto scaling is used to add additional instances as traffic increases under normal load the application runs 2 instances in the Auto Scaling group but at peak it can scale 3x in size. The application instances need to communicate with the payment service over the Internet which requires whitelisting of all public IP addresses used to communicate with it. A maximum of 4 whitelisting IP addresses are allowed at a time and can be added through an API.
How should they architect their solution?

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

Answer: D

NEW QUESTION 158
You are running a news website in the eu-west-1 region that updates every 15 minutes. The website has a world-wide audience it uses an Auto Scaling group behind an Elastic Load Balancer and an Amazon
RDS database Static content resides on Amazon S3, and is distributed through Amazon CIoudFront. Your Auto Scaling group is set to trigger a scale up event at 60% CPU utilization, you use an Amazon RDSextra large DB instance with 10.000 Provisioned IOPS its CPU utilization is around 80%. While freeable memory is in the 2 GB range.
Web analytics reports show that the average load time of your web pages is around 1.5 to 2 seconds, but your SEO consultant wants to bring down the average load time to under 0.5 seconds.
How would you improve page load times for your users? (Choose 3 answers)

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

Answer: ABD

NEW QUESTION 161

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 163
Your company previously configured a heavily used, dynamically routed VPN connection between your on-premises data center and AWS. You recently provisioned a DirectConnect connection and would like to start using the new connection. After configuring DirectConnect settings in the AWS Console, which of the following options win provide the most seamless transition for your users?

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

Answer: D

NEW QUESTION 164
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 167
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 sewers, two database sewers and one NAT instance tor a total of seven EC2 instances The web. Application and database sewers 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 (raffle 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 caused? (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 171
Your company produces customer commissioned one-of-a-kind skiing helmets combining nigh fashion with custom technical enhancements Customers can show off their IndMduality 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 manages 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 manages 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 auto-scaling group of C3 with SR-IOV (Single Root I/O virtualization).

Answer: B

NEW QUESTION 173
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 177
A company is building a voting system for a popular TV show, viewers win watch the performances then visit the show's website to vote for their favorite performer. It is expected that in a short period of time after the show has finished the site will receive millions of visitors. The visitors will first login to the site using their Amazon.com credentials and then submit their vote. After the voting is completed the page will display the vote totals. The company needs to build the site such that can handle the rapid influx of traffic while maintaining good performance but also wants to keep costs to a minimum. Which of the design patterns below should they use?

A. Use CIoudFront and an Elastic Load balancer in front of an auto-scaled set of web servers, the web servers will first call the Login With Amazon service to authenticate the user then process the users vote and store the result into a multi-AZ Relational Database Service instance.
B. Use CIoudFront and the static website hosting feature of S3 with the Javascript SDK to call the Login With Amazon service to authenticate the user, use IAM Roles to gain permissions to a DynamoDB tableto store the users vote.
C. Use CIoudFront and an Elastic Load Balancer in front of an auto-scaled set of web servers, the web servers will first call the Login with Amazon service to authenticate the user, the web servers will process the users vote and store the result into a DynamoDB table using IAM Roles for EC2 instances to gain permissions to the DynamoDB table.
D. Use CIoudFront and an Elastic Load Balancer in front of an auto-scaled set of web servers, the web sewers win process the users vote and store the result into an SQS queue using IAM Roles for EC2 Instances to gain permissions to the SQS queu
E. A set of application sewers will then retrieve the items from the queue and store the result into a DynamoDB table.

Answer: D

NEW QUESTION 179
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 183
Which is a valid Amazon Resource name (ARN) for IAM?

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

Answer: B

NEW QUESTION 185
Your fortune 500 company has under taken a TCO analysis evaluating the use of Amazon S3 versus acquiring more hardware The outcome was that ail employees would be granted access to use Amazon S3 for storage of their personal documents.

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

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

Answer: ABD

NEW QUESTION 186
You've been brought in as solutions architect to assist an enterprise customer with their migration of an e-commerce platform to Amazon Virtual Private Cloud (VPC) The previous architect has already deployed a 3-tier VPC.
The configuration is as follows: VPC: vpc-2f8bc447
IGW: igw-2d8bc445 NACL: ad-208bc448
Subnets and Route Tables: Web sewers: subnet-258bc44d
Application servers: subnet-248bc44c Database sewers: subnet-9189c6f9 Route Tables:
rrb-218bc449 rtb-238bc44b Associations:
subnet-258bc44d : rtb-218bc449 subnet-248bc44c : rtb-238bc44b subnet-9189c6f9 : rtb-238bc44b
You are now ready to begin deploying EC2 instances into the VPC Web servers must have direct access to the internet Application and database servers cannot have direct access to the internet.
Which configuration below will allow you the ability to remotely administer your application and database servers, as well as allow these sewers to retrieve updates from the Internet?

A. Create a bastion and NAT instance in subnet-258bc44d, and add a route from rtb- 238bc44b to the NAT instance.
B. Add a route from rtb-238bc44b to igw-2d8bc445 and add a bastion and NAT instance within subnet-248bc44c.
C. Create a bastion and NAT instance in subnet-248bc44c, and add a route from rtb- 238bc44b to subneb258bc44d.
D. Create a bastion and NAT instance in subnet-258bc44d, add a route from rtb-238bc44b toIgw-2d8bc445, and a new NACL that allows access between subnet-258bc44d and subnet-248bc44

Answer: A

NEW QUESTION 187
An administrator is using Amazon CIoudFormation 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 CIoudFormation 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 CIoudFormation 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 188
You have an application running on an EC2 Instance which will allow users to download flies 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 IANI 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 bucke
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 bucke
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 189
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 194
After launching an instance that you intend to serve as a NAT (Network Address Translation) device in a public subnet you modify your route tables to have the NAT device be the target of internet bound traffic of your private subnet. When you try and make an outbound connection to the internet from an instance in the private subnet, you are not successful. Which of the following steps could resolve the issue?

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

Answer: A

NEW QUESTION 198
Your application provides data transformation services. Files containing data to be transformed are first uploaded to Amazon S3 and then transformed by a fileet 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 leve
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 priorit
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 queu
G. Each message contains the priority leve
H. Transformation instances poll high-priority messages first.

Answer: C

NEW QUESTION 199
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 202
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 205
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 208
An ERP application is deployed across multiple AZs in a single region. In the event of failure, the Recovery Time Objective (RTO) must be less than 3 hours, and

the Recovery Point Objective (RPO) must be 15 minutes the customer realizes that data corruption occurred roughly 1.5 hours ago.
What DR strategy could be used to achieve this RTO and RPO in the event of this kind of failure?

A. Take hourly DB backups to S3, with transaction logs stored in S3 every 5 minutes.
B. Use synchronous database master-slave replication between two availability zones.
C. Take hourly DB backups to EC2 Instance store volumes with transaction logs stored In S3 every 5 minutes.
D. Take 15 minute DB backups stored In Glacier with transaction logs stored in S3 every 5 minute

Answer: A

NEW QUESTION 209
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 gatewa
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 gatewa
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: on to the Internet via an Internet gateway, the other to the on-premises network via the VPN gatewa
H. Use this routing table across all subnets in the VPC.
I. Configure two routing tables: on that has a default router via the Internet gateway, and other that has a default route via the VPN gatewa
J. Associate both routing tables with each VPC subnet.

Answer: A

NEW QUESTION 212
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 217
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 221
How is AWS readily distinguished from other vendors in the traditional IT computing landscape?

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

R. Dynami
S. Secur
T. Experience

Answer: C

NEW QUESTION 224
Your customer is willing to consolidate their log streams (access logs application logs security logs etc.) in one single system. Once consolidated, the customer wants to analyze these logs in real time based on heuristics. From time to time, the customer needs to validate heuristics, which requires going back to data samples extracted from the last 12 hours?
What is the best approach to meet your customer's requirements?

A. Send all the log events to Amazon SQS, setup an Auto Scaling group of EC2 sewers to consume the logs and apply the heuristics.
B. Send all the log events to Amazon Kinesis, develop a client process to apply heuristics on the logs
C. Configure Amazon CIoudTraiI to receive custom logs, use EMR to apply heuristics the logs
D. Setup an Auto Scaling group of EC2 syslogd servers, store the logs on S3, use EMR to apply heuristics on the logs

Answer: B

NEW QUESTION 228
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Certified-Solutions-Architect-Professional Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Certified-Solutions-Architect-Professional Product From:

https://www.2passeasy.com/dumps/AWS-Certified-Solutions-Architect-Professional/

Money Back Guarantee

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