AWS-Certified-Solutions-Architect-Professional Dumps

Amazon AWS Certified Solutions Architect Professional

https://www.certleader.com/AWS-Certified-Solutions-Architect-Professional-dumps.html

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

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

Answer: C

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

NEW QUESTION 2
An IAM user is trying to perform an action on an object belonging to some other root account's bucket. Which of the below mentioned options will AWS S3 not verify?

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

Answer: B

Explanation:
If the IAM user is trying to perform some action on the object belonging to another AWS user's bucket, S3 will verify whether the owner of the IAM user has given sufficient permission to him. It also verifies the policy for the bucket as well as the policy defined by the object owner.
Reference:
http://docs.aws.amazon.com/AmazonS3/latest/dev/access-control-auth-workflow-object-operation.html

NEW QUESTION 3
An organization is setting a website on the AWS VPC. The organization has blocked a few IPs to avoid a D-DOS attack. How can the organization configure that a request from the above mentioned IPs does not access the application instances?

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

Answer: D

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

NEW QUESTION 4
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.html

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

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

Answer: D

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

NEW QUESTION 6
In Amazon IAM, what is the maximum length for a role name?

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

Answer: C

Explanation:
In Amazon IAM, the maximum length for a role name is 64 characters.
Reference: http://docs.aws.amazon.com/IANI/latest/UserGuide/LimitationsOnEntities.html

NEW QUESTION 7
A user is planning to host a web server as well as an app server on a single EC2 instance which is a part of the public subnet of a VPC. How can the user setup to have two separate public IPs and separate security groups for both the application as well as the web server?

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

Answer: B

Explanation:
If you need to host multiple websites(with different IPs) on a single EC2 instance, the following is the suggested method from AWS.
Launch a VPC instance with two network interfaces
Assign elastic IPs from VPC EIP pool to those interfaces (Because, when the user has attached more than one network interface with an instance, AWS cannot assign public IPs to them.)
Assign separate Security Groups if separate Security Groups are needed
This scenario also helps for operating network appliances, such as firewalls or load balancers that have multiple private IP addresses for each network interface.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Mu|tip|e|P.htmI

NEW QUESTION 8
How many g2.2xIarge on-demand instances can a user run in one region without taking any limit increase approval from AWS?

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

Answer: C

Explanation:
Generally AWS EC2 allows running 20 on-demand instances and 100 spot instances at a time. This limit can be increased by requesting at
https://aws.amazon.com/contact-us/ec2-request. Excluding certain types of instances, the limit is lower than mentioned above. For g2.2xlarge, the user can run only 5
on-demand instance at a time.
Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_|imits.html#limits_ec2

NEW QUESTION 9
A user has created a NIySQL RDS instance with PIOPS. Which of the below mentioned statements will help user understand the advantage of PIOPS?

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

Answer: C

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

NEW QUESTION 10
In Amazon ElastiCache, the failure of a single cache node can have an impact on the availability of your application and the load on your back-end database while ElastiCache provisions a replacement for the failed cache node and it get repopulated. Which of the following is a solution to reduce this potential availability impact?

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

Answer: B

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

NEW QUESTION 10
What bandwidths do AWS Direct Connect currently support?

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

Answer: D

Explanation:
AWS Direct Connection currently supports 1Gbps and 10 Gbps.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

NEW QUESTION 14
The Principal element of an IAM policy refers to the specific entity that should be allowed or denied permission, whereas the translates to everyone except the specified entity.

A. NotPrincipa|
B. Vendor
C. Principal
D. Action

Answer: A

Explanation:
The element NotPrincipa| that is included within your IAM policy statements allows you to specify an exception to a list of principals to whom the access to a specific resource is either allowed or denied. Use the NotPrincipal element to specify an exception to a list of principals. For example, you can deny access to all principals except the one named in the NotPrincipa| element.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html#Principal

NEW QUESTION 15
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 19
A user has configured EBS volume with PIOPS. The user is not experiencing the optimal throughput. Which of the following could not be factor affecting I/O performance of that EBS volume?

A. EBS bandwidth of dedicated instance exceeding the PIOPS
B. EBS volume size

C. EC2 bandwidth
D. Instance type is not EBS optimized

Answer: B

Explanation:
If the user is not experiencing the expected IOPS or throughput that is provisioned, ensure that the EC2 bandwidth is not the limiting factor, the instance is EBS-optimized (or include 10 Gigabit network connectMty) and the instance type EBS dedicated bandwidth exceeds the IOPS more than he has provisioned.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-io-characteristics.html

NEW QUESTION 24
IAM Secure And Scalable is an organization which provides scalable and secure SAAS to its clients. They are planning to host a web server and App server on AWS VPC as separate tiers. The organization wants to implement the scalability by configuring Auto Scaling and load balancer with their app servers (middle tier) too. Which of the below mentioned options suits their requirements?

A. Since ELB is internet facing, it is recommended to setup HAProxy as the Load balancer within the VPC.
B. Create an Internet facing ELB with VPC and configure all the App servers with it.
C. The user should make ELB with EC2-CLASSIC and enable SSH with it for security.
D. Create an Internal Load balancer with VPC and register all the App sewers with i

Answer: D

Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances.
There are two ELBs available with VPC: internet facing and internal (private) ELB. For internal servers, such as App sewers the organization can create an internal load balancer in their VPC and then place back-end application instances behind the internal load balancer. The internal load balancer will route requests to the back-end application instances, which are also using private IP addresses and only accept requests from the internal load balancer.
Reference:
http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveIoperGuide/vpc-Ioadbalancer-types.html

NEW QUESTION 26
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 31
By default, what is the maximum number of Cache Nodes you can run in Amazon EIastiCache?

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

Answer: A

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

NEW QUESTION 34
Does an AWS Direct Connect location provide access to Amazon Web Services in the region it is associated with as well as access to other US regions?

A. No, it provides access only to the region it is associated with.
B. No, it provides access only to the US regions other than the region it is associated with.
C. Yes, it provides access.
D. Yes, it provides access but only when there's just one Availability Zone in the regio

Answer: C

Explanation:
An AWS Direct Connect location provides access to Amazon Web Services in the region it is associated with, as well as access to other US regions. For example,

you can provision a single connection to any AWS Direct Connect location in the US and use it to access public AWS services in all US Regions and AWS GovCloud (US).
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

NEW QUESTION 35
The CFO of a company wants to allow one of his employees to view only the AWS usage report page. Which of the below mentioned IAM policy statements allows the user to have access to the AWS usage report page?

A. "Effect": "AIIow", "Action": ["Describe"], "Resource": "BiIIing"
B. "Effect": "AIIow", "Action": ["aws-portal: ViewBi||ing"], "Resource": "*"
C. "Effect": "AIIow", "Action": ["aws-portal:ViewUsage"], "Resource": "*"
D. "Effect": "AIIow", "Action": ["AccountUsage], "Resource": "*"

Answer: C

Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the CFO wants to allow only AWS usage report page access, the policy for that IAM user will be as given below:
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "A||ow", "Action": [
"aws-portaI:ViewUsage"
]!
"Resource": "*"
} I
}
Reference: http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/billing-permissions-ref.html

NEW QUESTION 38
The user has provisioned the PIOPS volume with an EBS optimized instance. Generally speaking, in which I/O chunk should the bandwidth experienced by the user be measured by AWS?

A. 128 KB
B. 256 KB
C. 64 KB
D. 32 KB

Answer: B

Explanation:
IOPS are input/output operations per second. Amazon EBS measures each I/O operation per second (that is 256 KB or smaller) as one IOPS.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-io-characteristics.html

NEW QUESTION 40
A user is planning to use EBS for his DB requirement. The user already has an EC2 instance running in the VPC private subnet. How can the user attach the EBS volume to a running instance?

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

Answer: A

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

NEW QUESTION 44
A user is trying to create a vault in AWS Glacier. The user wants to enable notifications. In which of the below mentioned options can the user enable the notifications from the AWS console?

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

Answer: D

Explanation:
From AWS console the user can configure to have notifications sent to Amazon Simple Notifications Service (SNS). The user can select specific jobs that, on completion, will trigger the notifications such as Vault Inventory Retrieval Job Complete and Archive Retrieval Job Complete.
Reference: http://docs.aws.amazon.com/amazonglacier/latest/dev/configuring-notifications-console.html

NEW QUESTION 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.htmI

NEW QUESTION 47
ExamKilIer has three separate departments and each department has their own AWS accounts. The HR department has created a file sharing site where all the on roll employees' data is uploaded. The Admin department uploads data about the employee presence in the office to their DB hosted in the VPC. The Finance department needs to access data from the HR department to know the on roll employees to calculate the salary based on the number of days that an employee is present in the office.
How can ExamKil|er setup this scenario?

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

Answer: C

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. A VPC peering connection allows the user to route traffic between the peer VPCs using private IP addresses as if they are a part of the same network. This is helpful when one VPC from the same or different AWS account wants to connect with resources of the other VPC.
Reference:
http://docs.aws.amazon.com/AmazonVPC/Iatest/PeeringGuide/peering-configurations-full-access.htmI#t hree-vpcs-full-access

NEW QUESTION 52
What is the maximum length for an instance profile name in AWS IAM?

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

Answer: B

Explanation:
The maximum length for an instance profile name is 128 characters.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html

NEW QUESTION 55
A user has configured EBS volume with PIOPS. The user is not experiencing the optimal throughput. Which of the following could not be factor affecting I/O performance of that EBS volume?

A. EBS bandwidth of dedicated instance exceeding the PIOPS
B. EC2 bandwidth
C. EBS volume size
D. Instance type is not EBS optimized

Answer: C

Explanation:
If the user is not experiencing the expected IOPS or throughput that is provisioned, ensure that the EC2 bandwidth is not the limiting factor, the instance is EBS-optimized (or include 10 Gigabit network
connectMty) and the instance type EBS dedicated bandwidth exceeds the IOPS more than he has provisioned.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-io-characteristics.html

NEW QUESTION 57
If a single condition within an IAM policy includes multiple values for one key, it will be evaluated using a logical .

A. OR
B. NAND

C. NOR
D. AND

Answer: A

Explanation:
If a single condition within an IAM policy includes multiple values for one key, it will be evaluated using a logical OR.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html

NEW QUESTION 61
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 64
In a VPC, can you modify a set of DHCP options after you create them?

A. Yes, you can modify a set of DHCP options within 48 hours after creation and there are no VPCs associated with them.
B. Yes, you can modify a set of DHCP options any time after you create them.
C. No, you can't modify a set of DHCP options after you create them.
D. Yes, you can modify a set of DHCP options within 24 hours after creatio

Answer: C

Explanation:
After you create a set of DHCP options, you can't modify them. If you want your VPC to use a different set of DHCP options, you must create a new set and associate them with your VPC. You can also set up your VPC to use no DHCP options at all.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_DHCP_Options.html

NEW QUESTION 66
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.html

NEW QUESTION 68
By default, temporary security credentials for an IAM user are valid for a maximum of 12 hours, but you can request a duration as long as hours.

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

Answer: B

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

NEW QUESTION 71
With Amazon Elastic MapReduce (Amazon EMR) you can analyze and process vast amounts of data. The cluster is managed using an open-source framework called Hadoop.
You have set up an application to run Hadoop jobs. The application reads data from DynamoDB and generates a temporary file of 100 TBs.
The whole process runs for 30 minutes and the output of the job is stored to S3. Which of the below mentioned options is the most cost effective solution in this

case?

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

Answer: B

Explanation:
AWS EC2 Spot Instances allow the user to quote his own price for the EC2 computing capacity. The user can simply bid on the spare Amazon EC2 instances and run them whenever his bid exceeds the current Spot Price. The Spot Instance pricing model complements the On-Demand and Reserved Instance pricing models, providing potentially the most cost-effective option for obtaining compute capacity, depending on the application. The only challenge with a Spot Instance is data persistence as the instance can be terminated whenever the spot price exceeds the bid price.
In the current scenario a Hadoop job is a temporary job and does not run for a longer period. It fetches data from a persistent DynamoDB. Thus, even if the instance gets terminated there will be no data loss and the job can be re-run. As the output files are large temporary files, it will be useful to store data on ephermal storage for cost savings.
Reference: http://aws.amazon.com/ec2/purchasing-options/spot-instances/

NEW QUESTION 76
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/AmazonElastiCache/latest/UserGuide/AutoDiscovery.HowAutoDiscoveryW orks.htmI

NEW QUESTION 79
An organization has created multiple components of a single application for compartmentalization. Currently all the components are hosted on a single EC2 instance. Due to security reasons the organization wants to implement two separate SSLs for the separate modules although it is already using VPC. How can the organization achieve this with a single instance?

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

Answer: B

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

NEW QUESTION 80
A user is thinking to use EBS PIOPS volume. Which of the below mentioned options is a right use case for the PIOPS EBS volume?

A. Analytics
B. System boot volume
C. Nlongo DB
D. Log processing

Answer: C

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. 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 business applications, database workloads, such as NoSQL DB, RDBMS, etc. Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVo|umeTypes.htm|

NEW QUESTION 83
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 88
Which of the following is true while using an IAM role to grant permissions to applications running on Amazon EC2 instances?

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

Answer: D

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

NEW QUESTION 91
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'll need to implement that functionality in a custom task runner.
Reference:
http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/dp-how-remote-taskrunner-client.html

NEW QUESTION 96
With respect to AWS Lambda permissions model, at the time you create a Lambda function, you specify an IAM role that AWS Lambda can assume to execute your Lambda function on your behalf. This role is also referred to as the role.

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

Answer: B

Explanation:
Regardless of how your Lambda function is invoked, AWS Lambda always executes the function. At the time you create a Lambda function, you specify an IAM role that AWS Lambda can assume to execute your Lambda function on your behalf. This role is also referred to as the execution role.
Reference: http://docs.aws.amazon.com/lambda/latest/dg/lambda-dg.pdf

NEW QUESTION 99
An organization is planning to use NoSQL DB for its scalable data needs. The organization wants to host an application securely in AWS VPC. What action can be recommended to the organization?

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

Answer: A

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

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Introduction.htm|

NEW QUESTION 103
What happens when Dedicated instances are launched into a VPC?

A. If you launch an instance into a VPC that has an instance tenancy of dedicated, you must manually create a Dedicated instance.
B. If you launch an instance into a VPC that has an instance tenancy of dedicated, your instance is created as a Dedicated instance, only based on the tenancy of the instance.
C. If you launch an instance into a VPC that has an instance tenancy of dedicated, your instance isautomatically a Dedicated instance, regardless of the tenancy of the instance.
D. None of these are tru

Answer: C

Explanation:
If you launch an instance into a VPC that has an instance tenancy of dedicated, your instance is automatically a Dedicated instance, regardless of the tenancy of the instance.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/dedicated-instance.html

NEW QUESTION 108
An organization is setting up RDS for their applications. The organization wants to secure RDS access with VPC. Which of the following options is not required while designing the RDS with VPC?

A. The organization must create a subnet group with public and private subnet
B. Both the subnets can be in the same or separate AZ.
C. The organization should keep minimum of one IP address in each subnet reserved for RDS failover.
D. If the organization is connecting RDS from the internet it must enable the VPC attributes DNS hostnames and DNS resolution.
E. The organization must create a subnet group with VPC using more than one subnet which are a part of separate AZs.

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, such as RDS into a virtual network that the user has defined. Subnets are segments of a VPC's IP address range that the user can designate to a group of VPC resources based on security and operational needs. A DB subnet group is a collection of subnets (generally private) that the user can create in a VPC and assign to the RDS DB instances. A DB subnet group allows the user to specify a particular VPC when creating the DB instances.
Each DB subnet group should have subnets in at least two Availability Zones in a given region. If the RDS instance is required to be accessible from the internet the organization must enable the VPC attributes, DNS hostnames and DNS resolution. For each RDS DB instance that the user runs in a VPC, he should reserve at least one address in each subnet in the DB subnet group for use by Amazon RDS for recovery actions.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.html

NEW QUESTION 109
You create a VPN connection, and your VPN device supports Border Gateway Protocol (BGP). Which of the following should be specified to configure the VPN connection?

A. Classless routing
B. Classfull routing
C. Dynamic routing
D. Static routing

Answer: C

Explanation:
If you create a VPN connection, you must specify the type of routing that you plan to use, which will depend upon on the make and model of your VPN devices. If your VPN device supports Border Gateway Protocol (BGP), you need to specify dynamic routing when you configure your VPN connection. If your device does not support BGP, you should specify static routing.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.htmI

NEW QUESTION 112
An organization has developed an application which provides a smarter shopping experience. They need to show a demonstration to various stakeholders who may not be able to access the in premise
application so they decide to host a demo version of the application on AWS. Consequently they will need a fixed elastic IP attached automatically to the instance when it is launched.
In this scenario which of the below mentioned options will not help assign the elastic IP automatically?

A. Write a script which will fetch the instance metadata on system boot and assign the public IP using that metadata.
B. Provide an elastic IP in the user data and setup a bootstrapping script which will fetch that elastic IP and assign it to the instance.
C. Create a controlling application which launches the instance and assigns the elastic IP based on the parameter provided when that instance is booted.
D. Launch instance with VPC and assign an elastic IP to the primary network interfac

Answer: A

Explanation:
EC2 allows the user to launch On-Demand instances. If the organization is using an application temporarily only for demo purposes the best way to assign an elastic IP would be:
Launch an instance with a VPC and assign an EIP to the primary network interface. This way on every instance start it will have the same IP Create a
bootstrapping script and provide it some metadata, such as user data which can be used to assign an EIP Create a controller instance which can schedule the start and stop of the instance and provide an EIP as a parameter so that the controller instance can check the instance boot and assign an EIP
The instance metadata gives the current instance data, such as the public/private IP. It can be of no use for assigning an EIP.

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

NEW QUESTION 114
Can a Direct Connect link be connected directly to the Internet?

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

Answer: D

Explanation:
AWS Direct Connect is a network service that provides an alternative to using the Internet to utilize AWS cloud service. Hence, a Direct Connect link cannot be connected to the Internet directly.
Reference: http://aws.amazon.com/directconnect/faqs/

NEW QUESTION 117
ExamKiIIer 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. ExamKiIIer wants to setup a centralized server which can connect with the LMS of each tenant upgrade if required. ExamKiIIer 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.htmI# many-vpcs-full-acces

NEW QUESTION 118
A user has created a VPC with CIDR 20.0.0.0/16. The user has created one subnet with CIDR 20.0.0.0/16 in this VPC. The user is trying to create another subnet with the same VPC for CIDR 20.0.0.1/24. What will happen in this scenario?

A. The VPC will modify the first subnet CIDR automatically to allow the second subnet IP range
B. The second subnet will be created
C. It will throw a CIDR overlaps error
D. It is not possible to create a subnet with the same CIDR as VPC

Answer: C

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. The user can create a subnet with the same size of VPC. However, he cannot create any other subnet since the CIDR of the second subnet will conflict with the first subnet.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

NEW QUESTION 120
In Amazon Cognito what is a silent push notification?

A. It is a push message that is received by your application on a user's device that will not be seen by 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 123
AWS has launched T2 instances which come with CPU usage credit. An organization has a requirement which keeps an instance running for 24 hours. However, the organization has high usage only during 11 AM to 12 PM. The organization is planning to use a T2 small instance for this purpose.
If the organization already has multiple instances running since Jan 2012, which of the below mentioned options should the organization implement while

launching a T2 instance?

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

Answer: C

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

NEW QUESTION 125
How does AWS Data Pipeline execute actMties on on-premise resources or AWS resources that you manage?

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

Answer: A

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

NEW QUESTION 128
An organization is planning to host a web application in the AWS VPC. The organization does not want to host a database in the public cloud due to statutory requirements. How can the organization setup in this scenario?

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

Answer: A

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

NEW QUESTION 130
A user is trying to create a PIOPS EBS volume with 4000 IOPS and 100 GB size. AWS does not allow the user to create this volume. What is the possible root cause for this?

A. PIOPS is supported for EBS higher than 500 GB size
B. The maximum IOPS supported by EBS is 3000
C. The ratio between IOPS and the EBS volume is higher than 30
D. The ratio between IOPS and the EBS volume is lower than 50

Answer: C

Explanation:
A Provisioned IOPS (SSD) volume can range in size from 4 GiB to 16 TiB and you can provision up to 20,000 IOPS per volume. The ratio of IOPS provisioned to the volume size requested should be a maximum of 30; for example, a volume with 3000 IOPS must be atleast 100 GB.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVo|umeTypes.htmI#EBSVolumeTypes_pio ps

NEW QUESTION 133
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 137
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 140
What is the network performance offered by the c4.8xlarge instance in Amazon EC2?

A. Very High but variable
B. 20 Gigabit
C. 5 Gigabit
D. 10 Gigabit

Answer: D

Explanation:
Networking performance offered by the c4.8xlarge instance is 10 Gigabit. Reference: http://aws.amazon.com/ec2/instance-types/

NEW QUESTION 145
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 146
Do you need to use Amazon Cognito to use the Amazon Mobile Analytics service?

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

Answer: A

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

NEW QUESTION 148
You want to use Amazon Redshift and you are planning to deploy dw1.8xlarge nodes. What is the minimum amount of nodes that you need to deploy with this kind

of configuration?

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

Answer: D

Explanation:
For a single-node configuration in Amazon Redshift, the only option available is the smallest of the two options. The 8XL extra-large nodes are only available in a multi-node configuration
Reference: http://docs.aws.amazon.com/redshift/latest/mgmt/working-with-c|usters.htmI

NEW QUESTION 149
Mike is appointed as Cloud Consultant in ExamKi|Ier.com. ExamKil|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
ExamKilIer.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 ExamKil|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 150
A user is hosting a public website on AWS. The user wants to have the database and the app server on the AWS VPC. The user wants to setup a database that can connect to the Internet for any patch upgrade but cannot receive any request from the internet. How can the user set this up?

A. Setup DB in a private subnet with the security group allowing only outbound traffic.
B. Setup DB in a public subnet with the security group allowing only inbound data.
C. Setup DB in a local data center and use a private gateway to connect the application with DB.
D. Setup DB in a private subnet which is connected to the internet via NAT for outbound.

Answer: D

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. AWS provides two features that the user can use to increase security in VPC: security groups and network ACLs. When the user wants to setup both the DB and App on VPC, the user should make one public and one private subnet. The DB should be hosted in a private subnet and instances in that subnet cannot reach the internet. The user can allow an instance in his VPC to initiate outbound connections to the internet but prevent unsolicited inbound connections from the internet by using a Network Address Translation (NAT) instance.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

NEW QUESTION 155
In Amazon ElastiCache, the default cache port is:

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

Answer: D

Explanation:
In Amazon ElastiCache, you can specify a new port number for your cache cluster, which by default is 11211 for Memcached and 6379 for Redis.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/Iatest/UserGuide/GettingStarted.AuthorizeAccess.html|

NEW QUESTION 156
Is there any way to own a direct connection to Amazon Web Services?

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

Answer: D

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

NEW QUESTION 157
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/examkiIler/
B. https://signin.aws.amazon.com/examki|Ier/
C. https://examkiller.signin.aws.amazon.com/999988887777/console/
D. https://999988887777.signin.aws.amazon.com/console/

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/console/. 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 159
Your company has recently extended its datacenter into a VPC on AVVS to add burst computing capacity as needed Members of your Network Operations Center need to be able to go to the AWS Management Console and administer Amazon EC2 instances as necessary You don't want to create new IAM users for each NOC member and make those users sign in again to the AWS Management Console Which option below will meet the needs for your NOC members?

A. Use OAuth 2 0 to retrieve temporary AWS security credentials to enable your NOC members to sign in to the AWS Management Console.
B. Use web Identity Federation to retrieve AWS temporary security credentials to enable your NOC members to sign in to the AWS IV|anagement Console.
C. Use your on-premises SAML 2.0-compliant identity provider (IDP) to grant the NOC members federated access to the AWS Management Console via the AWS single sign-on (SSO) endpoint.
D. Use your on-premises SAML 2.0-compliam identity provider (IDP) to retrieve temporary security credentials to enable NOC members to sign in to the AWS Management Console.

Answer: D

NEW QUESTION 161
Your startup wants to implement an order fulfillment process for selling a personalized gadget that needs an average of 3-4 days to produce with some orders taking up to 6 months you expect 10 orders per day on your first day. 1000 orders per day after 6 months and 10,000 orders after 12 months.
Orders coming in are checked for consistency men dispatched to your manufacturing plant for production quality control packaging shipment and payment processing If the product does not meet the quality standards at any stage of the process employees may force the process to repeat a step Customers are notified via email about order status and any critical issues with their orders such as payment failure.
Your case architecture includes AWS Elastic Beanstalk for your website with an RDS MySQL instance for customer data and orders.
How can you implement the order fulfillment process while making sure that the emails are delivered reliably?

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

Answer: C

NEW QUESTION 165
You are tasked with moving a legacy application from a virtual machine running Inside your datacenter to an Amazon VPC Unfortunately this app requires access to a number of on-premises services and no one who configured the app still works for your company. Even worse there's no documentation for it.
What will allow the application running inside the VPC to reach back and access its internal dependencies
without being reconfigured? (Choose 3 answers)

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

Answer: ADF

NEW QUESTION 170
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 175
You currently operate a web application In the AWS US-East region The application runs on an
auto-scaled layer of EC2 instances and an RDS Multi-AZ database Your IT security compliance officer has tasked you to develop a reliable and durable logging
solution to track changes made to your EC2.IAM And RDS resources. The solution must ensure the integrity and confidentiality of your log data. Which of these
solutions would you recommend?

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

Answer: A

NEW QUESTION 176
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 177
An AWS customer is deploying an application mat is composed of an AutoScaling group of EC2 Instances.
The customers security policy requires that every outbound connection from these instances to any other service within the customers Virtual Private Cloud must
be authenticated using a unique x 509 certificate that contains the specific instance-id.
In addition an x 509 certificates must Designed by the customer's Key management service in order to be trusted for authentication.
Which of the following configurations will support these requirements?

A. Configure an IAM Role that grants access to an Amazon S3 object containing a signed certificate and configure me Auto Scaling group to launch instances with
this role Have the instances bootstrap get the certificate from Amazon S3 upon first boot.
B. Embed a certificate into the Amazon Machine Image that is used by the Auto Scaling group Have the launched instances generate a certificate signature
request with the instance's assigned instance-id to the Key management service for signature.
C. Configure the Auto Scaling group to send an SNS notification of the launch of a new instance to the trusted key management servic
D. Have the Key management service generate a signed certificate and send it directly to the newly launched instance.
E. Configure the launched instances to generate a new certificate upon first boot Have the Key management service poll the Auto Scaling group for associated
instances and send new instances a certificate signature (hat contains the specific instance-id.

Answer: A

NEW QUESTION 181
Your company runs a customer facing event registration site This site is built with a 3-tier architecture with web and application tier servers and a MySQL database
The application requires 6 web tier sewers and 6 application tier servers for normal operation, but can run on a minimum of 65% server capacity and a single
NIySQL database. When deploying this application in a region with three availability zones (AZs) which architecture provides high availability?

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

Answer: D

NEW QUESTION 183
Your customer wishes to deploy an enterprise application to AWS which will consist of several web servers, several application servers and a small (50GB) Oracle database information is stored, both in the database and the file systems of the various servers. The backup system must support database recovery whole server and whole disk restores, and indMdual file restores with a recovery time of no more than two hours. They have chosen to use RDS Oracle as the database
Which backup architecture will meet these requirements?

A. Backup RDS using automated daily DB backups Backup the EC2 instances using AMIs and supplement with file-level backup to S3 using traditional enterprise backup software to provide file level restore
B. Backup RDS using a Multi-AZ Deployment Backup the EC2 instances using Amis, and supplement by copying file system data to S3 to provide file level restore.
C. Backup RDS using automated daily DB backups Backup the EC2 instances using EBS snapshots and supplement with file-level backups to Amazon Glacier using traditional enterprise backup software to provide file level restore
D. Backup RDS database to S3 using Oracle RMAN Backup the EC2 instances using Amis, and supplement with EBS snapshots for indMdual volume restore.

Answer: A

NEW QUESTION 186
Your company has HQ in Tokyo and branch offices all over the world and is using a logistics software with a multi-regional deployment on AWS in Japan, Europe and US

A. The logistic software has a 3-tierarchitecture and currently uses MySQL 5.6 for data persistenc
B. Each region has deployed its own database In the HQ region you run an hourly batch process reading data from every region to compute cross-regional reports that are sent by email to all offices this batch process must be completed as fast as possible to quickly optimize logistics how do you build the database architecture in order to meet the requirements'?
C. For each regional deployment, use RDS MySQL with a master in the region and a read replica in theHQ region
D. For each regional deployment, use NIySQL on EC2 with a master in the region and send hourly EBS snapshots to the HQ region
E. For each regional deployment, use RDS MySQL with a master in the region and send hourly RDS snapshots to the HQ region
F. For each regional deployment, use MySQL on EC2 with a master in the region and use S3 to copy data files hourly to the HQ region
G. Use Direct Connect to connect all regional MySQL deployments to the HQ region and reduce network latency for the batch process

Answer: A

NEW QUESTION 190
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 192
An enterprise wants to use a third-party SaaS application. The SaaS application needs to have access to issue several API commands to discover Amazon EC2 resources running within the enterprise's account The enterprise has internal security policies that require any outside access to their environment must conform to the principles of least prMlege and there must be controls in place to ensure that the credentials used by the SaaS vendor cannot be used by any other third party.
Which of the following would meet all of these conditions?

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

Answer: C

NEW QUESTION 195
You are responsible for a legacy web application whose server environment is approaching end of life You would like to migrate this application to AWS as quickly as possible, since the application environment currently has the following limitations:
The VM's single 10GB VNIDK is almost full; Nle virtual network interface still uses the 10IV|bps driver, which leaves your 100Mbps WAN connection completely underutilized;
It is currently running on a highly customized. Windows VM within a VMware environment; You do not have me installation media;
This is a mission critical application with an RTO (Recovery Time Objective) of 8 hours. RPO (Recovery Point Objective) of 1 hour. How could you best migrate this application to AWS while meeting your business continuity requirements?

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

Answer: A

NEW QUESTION 199
You are implementing a URL whitelisting system for a company that wants to restrict outbound HTTP'S connections to specific domains from their EC2-hosted applications you deploy a single EC2 instance running proxy software and configure It to accept traffic from all subnets and EC2 instances in the VPC. You configure the proxy to only pass through traffic to domains that you define in its whitelist configuration You have a nightly maintenance window or 10 minutes where ail instances fetch new software updates. Each update Is about 200MB In size and there are 500 instances In the VPC that routinely fetch updates After a few days you notice that some machines are failing to successfully download some, but not all of their updates within the maintenance window. The download URLs used for these updates are correctly listed in the proxy's whitelist configuration and you are able to access them manually using a web browser on the instances. What might be happening? (Choose 2 answers)

A. You are running the proxy on an undersized EC2 instance type so network throughput is not sufficient for all instances to download their updates in time.
B. You are running the proxy on a sufficiently-sized EC2 instance in a private subnet andits network throughput is being throttled by a NAT running on an undersized EC2 instance.
C. The route table for the subnets containing the affected EC2 instances is not configured to direct network traffic for the software update locations to the proxy.
D. You have not allocated enough storage to the EC2 instance running the proxy so the network buffer is filling up, causing some requests to fail.
E. You are running the proxy in a public subnet but have not allocated enough EIPs to support the needed network throughput through the Internet Gateway (IGW).

Answer: AB

NEW QUESTION 201
Company B is launching a new game app for mobile devices. Users will log into the game using their existing social media account to streamline data capture. Company B would like to directly save player data and scoring information from the mobile app to a DynamoDS table named Score Data When a user saves their game the progress data will be stored to the Game state S3 bucket. What is the best approach for storing data to DynamoDB and S3?

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

Answer: B

NEW QUESTION 203
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 206
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 211
Your website is serving on-demand training videos to your workforce. Videos are uploaded monthly in high resolution MP4 format. Your workforce is distributed globally often on the move and using company-provided tablets that require the HTTP Live Streaming (HLS) protocol to watch a video. Your company has no video transcoding expertise and it required you may need to pay for a consultant.

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

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

Answer: C

NEW QUESTION 213
Your website is serving on-demand training videos to your workforce. Videos are uploaded monthly in high resolution MP4 format. Your workforce is distributed
globally often on the move and using company-provided tablets that require the HTTP Live Streaming (HLS) protocol to watch a video. Your company has no video
transcoding expertise and it required you may need to pay for a consultant.
How do you implement the most cost-efficient architecture without compromising high availability and quality of video delivery'?

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

Answer: C

NEW QUESTION 216
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 218
You are running a successful multitier web application on AWS and your marketing department has asked you to add a reporting tier to the application. The
reporting tier will aggregate and publish status reports every 30 minutes from user-generated information that is being stored in your web application s database.
You are currently running a Multi-AZ RDS MySQL instance for the database tier. You also have implemented Elasticache as a database caching layer between the
application tier and database tier. Please select the answer that will allow you to successfully implement the reporting tier with as little impact as possible to your
database.

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

Answer: C

NEW QUESTION 223
Your system recently experienced down time during the troubleshooting process. You found that a new administrator mistakenly terminated several production
EC2 instances.
Which of the following strategies will help prevent a similar situation in the future? The administrator still must be able to:
launch, start stop, and terminate development resources. launch and start production instances.

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

Answer: B

NEW QUESTION 225
An administrator is using Amazon CloudFormation to deploy a three tier web application that consists of a web tier and application tier that will utilize Amazon DynamoDB for storage when creating the CloudFormation template which of the following would allow the application instance access to the DynamoDB tables without exposing API credentials?

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

Answer: C

NEW QUESTION 229
Your company has recently extended its datacenter into a VPC on AWS to add burst computing capacity as needed Members of your Network Operations Center need to be able to go to the AWS Management Console and administer Amazon EC2 instances as necessary You don't want to create new IAM users for each NOC member and make those users sign in again to the AWS Management Console Which option below will meet the needs for your NOC members?

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

Answer: D

NEW QUESTION 231
You are developing a new mobile application and are considering storing user preferences in AWS.2w This would provide a more uniform cross-device experience to users using multiple mobile devices to access the application. The preference data for each user is estimated to be 50KB in size Additionally 5 million customers are expected to use the application on a regular basis. The solution needs to be
cost-effective, highly available, scalable and secure, how would you design a solution to meet the above requirements?

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

Answer: B

NEW QUESTION 234
You deployed your company website using Elastic Beanstalk and you enabled log file rotation to S3. An Elastic Map Reduce job is periodically analyzing the logs on S3 to build a usage dashboard that you share with your CIO.
You recently improved overall performance of the website using Cloud Front for dynamic content delivery and your website as the origin.
After this architectural change, the usage dashboard shows that the traffic on your website dropped by an order of magnitude. How do you fix your usage dashboard'?

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

Answer: D

NEW QUESTION 235
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 237
Select the correct set of options. These are the initial settings for the default security group:

A. Allow no inbound traffic, Allow all outbound traffic and Allow instances associated with this security group to talk to each other
B. Allow all inbound traffic, Allow no outbound traffic and Allow instances associated with this security group to talk to each other
C. Allow no inbound traffic, Allow all outbound traffic and Does NOT allow instances associated with this security group to talk to each other
D. Allow all inbound traffic, Allow all outbound traffic and Does NOT allow instances associated with this security group to talk to each other

Answer: A

NEW QUESTION 240
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 243
When you put objects in Amazon S3, what is the indication that an object was successfully stored?

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

Answer: A

NEW QUESTION 244
A web company is looking to implement an intrusion detection and prevention system into their deployed VPC. This platform should have the ability to scale to thousands of instances running inside of the VPC. How should they architect their solution to achieve these goals?

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

Answer: C

NEW QUESTION 246
A company is storing data on Amazon Simple Storage Service (S3). The company's security policy mandates that data is encrypted at rest. Which of the following methods can achieve this?
Choose 3 answers

A. Use Amazon S3 server-side encryption with AWS Key Management Service managed keys.
B. Use Amazon S3 server-side encryption with customer-provided keys.
C. Use Amazon S3 server-side encryption with EC2 key pair.
D. Use Amazon S3 bucket policies to restrict access to the data at rest.
E. Encrypt the data on the client-side before ingesting to Amazon S3 using their own master key.
F. Use SSL to encrypt the data while in transit to Amazon S3.

Answer: ABE

NEW QUESTION 248
Your company is in the process of developing a next generation pet collar that collects biometric information to assist families with promoting healthy lifestyles for their pets Each collar will push 30kb of biometric data In JSON format every 2 seconds to a collection platform that will process and analyze the data providing health trending information back to the pet owners and veterinarians via a web portal Nlanagement has tasked you to architect the collection platform ensuring the following requirements are met.
Provide the ability for real-time analytics of the inbound biometric data Ensure processing of the biometric data is highly durable. Elastic and parallel The results of the analytic processing should be persisted for data mining
Which architecture outlined below win meet the initial requirements for the collection platform?

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

Answer: B

NEW QUESTION 251
The AWS IT infrastructure that AWS provides, complies with the following IT security standards, including:

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

Answer: ABC

NEW QUESTION 252
What does elasticity mean to AWS?

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

Answer: B

NEW QUESTION 255
Your company is storing millions of sensitive transactions across thousands of 100-GB files that must be encrypted in transit and at rest. Analysts concurrently depend on subsets of files, which can consume up to 5 TB of space, to generate simulations that can be used to steer business decisions. You are required to design an AWS solution that can cost effectively accommodate the long-term storage and in-flight subsets of data.

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

Answer: D

NEW QUESTION 258
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 CIoudTrail 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 259
......

Thank You for Trying Our Product

* 100% Pass or Money Back

    All our products come with a 90-day Money Back Guarantee.

* One year free update

    You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

    We currently serve more than 30,000,000 customers.

* Shop Securely

    All transactions are protected by VeriSign!

100% Pass Your AWS-Certified-Solutions-Architect-Professional Exam with Our Prep Materials Via below:

https://www.certleader.com/AWS-Certified-Solutions-Architect-Professional-dumps.html