

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

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 2

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

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

Answer: C

Explanation:

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

NEW QUESTION 3

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

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

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 4

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

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

Answer: A

Explanation:

The maximum number of VPCs allowed per region is 5.

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

NEW QUESTION 5

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/MultiIP.html>

NEW QUESTION 6

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

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

Answer: C

Explanation:

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

NEW QUESTION 7

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

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

Answer: A

Explanation:

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

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

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

Reference: http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/USVPC_creating_basic_lb.html

NEW QUESTION 8

What bandwidths do AWS Direct Connect currently support?

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

Answer: D

Explanation:

AWS Direct Connection currently supports 1Gbps and 10 Gbps.

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

NEW QUESTION 9

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. NotPrincipal
- B. Vendor
- C. Principal
- D. Action

Answer: A

Explanation:

The element `NotPrincipal` 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 `NotPrincipal` element.

Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html#Principal

NEW QUESTION 10

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

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

Answer: A

Explanation:

In Amazon VPC, you can assign any Private IP address to your instance as long as it is: Part of the associated subnet's IP address range

Not reserved by Amazon for IP networking purposes Not currently assigned to another interface Reference: <http://aws.amazon.com/vpc/faqs/>

NEW QUESTION 10

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 13

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

- A. Create five VPCs with the public subnet for the app server and setup the VPN gateway for each VPN to connect them indMdually.
- B. Use the AWS VPN CloudHub to communicate with multiple VPN connections.
- C. Use the AWS CloudGateway to communicate with multiple VPN connections.
- D. It is not possible to connect different data centers from a single VPC.

Answer: B

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. If the organization has multiple VPN connections, he can provide secure communication between sites using the AWS VPN CloudHub.

The VPN CloudHub operates on a simple hub-and-spoke model that the user can use with or without a VPC. This design is suitable for customers with multiple branch offices and existing internet connections who would like to implement a convenient, potentially low-cost hub-and-spoke model for primary or backup connectMty between remote offices.

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

NEW QUESTION 18

One of your AWS Data Pipeline actMties has failed consequently and has entered a hard failure state after retrying thrice. You want to try it again. Is it possible to increase the number of automatic retries to more than thrice?

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

Answer: D

Explanation:

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

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

NEW QUESTION 20

True or False: In Amazon ElastiCache replication groups of Redis, for performance tuning reasons, you can change the roles of the cache nodes within the replication group, with the primary and one of the replicas exchanging roles.

- A. True, however, you get lower performance.
- B. FALSE
- C. TRUE
- D. False, you must recreate the replication group to improve performance tuning

Answer: C

Explanation:

In Amazon ElastiCache, a replication group is a collection of Redis Cache Clusters, with one primary read-write cluster and up to five secondary, read-only clusters, which are called read replicas. You can change the roles of the cache clusters within the replication group, with the primary cluster and one of the replicas exchanging roles. You might decide to do this for performance tuning reasons.

Reference: <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/Replication.Redis.Groups.html>

NEW QUESTION 22

How much memory does the cr1.8xlarge instance type provide?

- A. 224 GB
- B. 124 GB
- C. 184 GB
- D. 244 GB

Answer: D

Explanation:

The CR1 instances are part of the memory optimized instances. They offer lowest cost per GB RAM among all the AWS instance families. CR1 instances are part of the new generation of memory optimized instances, which can offer up to 244 GB RAM and run on faster CPUs (Intel Xeon E5-2670 with NUMA support) in comparison to the NI2 instances of the same family. They support cluster networking for bandwidth intensive applications. cr1.8xlarge is one of the largest instance types of the CR1 family, which can offer 244 GB RAM.

Reference: <http://aws.amazon.com/ec2/instance-types/>

NEW QUESTION 26

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 servers with it

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 servers the organization can create an internal load balancer in their VPC and then place back-end application instances behind the internal load balancer. The internal load balancer will route requests to the back-end application instances, which are also using private IP addresses and only accept requests from the internal load balancer.

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/vpc-loadbalancer-types.html>

NEW QUESTION 27

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

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

Answer: C

Explanation:

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

Reference: <https://aws.amazon.com/elasticache/>

NEW QUESTION 28

An organization is having an application which can start and stop an EC2 instance as per schedule. The organization needs the MAC address of the instance to be registered with its software. The instance is launched in EC2-CLASSIC. How can the organization update the MAC registration every time an instance is booted?

- A. The organization should write a boot strapping script which will get the MAC address from the instance metadata and use that script to register with the application.
- B. The organization should provide a MAC address as a part of the user data
- C. Thus, whenever the instance is booted the script assigns the fixed MAC address to that instance.
- D. The instance MAC address never change
- E. Thus, it is not required to register the MAC address every time.

F. AWS never provides a MAC address to an instance; instead the instance ID is used for identifying the instance for any software registration.

Answer: A

Explanation:

AWS provides an on demand, scalable infrastructure. AWS EC2 allows the user to launch On-Demand instances. AWS does not provide a fixed MAC address to the instances launched in EC2-CLASSIC. If the instance is launched as a part of EC2-VPC, it can have an ENI which can have a fixed MAC. However, with EC2-CLASSIC, every time the instance is started or stopped it will have a new MAC address.

To get this MAC, the organization can run a script on boot which can fetch the instance metadata and get the MAC address from that instance metadata. Once the MAC is received, the organization can register that MAC with the software.

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

NEW QUESTION 31

Does Amazon RDS API provide actions to modify DB instances inside a VPC and associate them with DB Security Groups?

- A. Yes, Amazon does this but only for MySQL RDS.
- B. Yes
- C. No
- D. Yes, Amazon does this but only for Oracle RD

Answer: B

Explanation:

You can use the action Modify DB Instance, available in the Amazon RDS API, to pass values for the parameters DB Instance Identifier and DB Security Groups specifying the instance ID and the DB Security Groups you want your instance to be part of.

Reference: http://docs.aws.amazon.com/AmazonRDS/latest/APIReference/API_ModifyDBInstance.html

NEW QUESTION 36

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

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

Answer: B

Explanation:

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

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

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

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

NEW QUESTION 40

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

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

Answer: B

Explanation:

A pipeline definition specifies the business logic of your data management.

Reference: <http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/what-is-datapipeline.html>

NEW QUESTION 42

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

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

Answer: D

Explanation:

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

Reference:

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

NEW QUESTION 43

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 46

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 48

An organization is undergoing a security audit. The auditor wants to view the AWS VPC configurations as the organization has hosted all the applications in the AWS VPC. The auditor is from a remote place and wants to have access to AWS to view all the VPC records.

How can the organization meet the expectations of the auditor without compromising on the security of their AWS infrastructure?

- A. The organization should not accept the request as sharing the credentials means compromising on security.
- B. Create an IAM role which will have read only access to all EC2 services including VPC and assign that role to the auditor.
- C. Create an IAM user who will have read only access to the AWS VPC and share those credentials with the auditor.
- D. The organization should create an IAM user with VPC full access but set a condition that will not allow to modify anything if the request is from any IP other than the organization's data center.

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 VPC also works with IAM and the organization can create IAM users who have access to various VPC services.

If an auditor wants to have access to the AWS VPC to verify the rules, the organization should be careful before sharing any data which can allow making updates to the AWS infrastructure. In this scenario it is recommended that the organization creates an IAM user who will have read only access to the VPC. Share the above mentioned credentials with the auditor as it cannot harm the organization. The sample policy is given below:

```
{
  "Effect": "Allow",
  "Action": [ "ec2:DescribeVpcs", "ec2:DescribeSubnets",
    "ec2:DescribeInternetGateways", "ec2:DescribeCustomerGateways", "ec2:DescribeVpnGateways", "ec2:DescribeVpnConnections", "ec2:DescribeRouteTables",
    "ec2:DescribeAddresses", "ec2:DescribeSecurityGroups", "ec2:DescribeNetworkAcls", "ec2:DescribeDhcpOptions", "ec2:DescribeTags", "ec2:DescribeInstances"
  ],
  "Resource": "*"
}
```

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

NEW QUESTION 51

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 59

An organization is planning to create a secure scalable application with AWS VPC and ELB. The organization has two instances already running and each instance has an ENI attached to it in addition to a primary network interface. The primary network interface and additional ENI both have an elastic IP attached to it.

If those instances are registered with ELB and the organization wants ELB to send data to a particular EIP of the instance, how can they achieve this?

- A. The organization should ensure that the IP which is required to receive the ELB traffic is attached to a primary network interface.
- B. It is not possible to attach an instance with two ENIs with ELB as it will give an IP conflict error.
- C. The organization should ensure that the IP which is required to receive the ELB traffic is attached to an additional ENI.
- D. It is not possible to send data to a particular IP as ELB will send to any one EIP

Answer: A

Explanation:

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 the internet facing ELB it is required that the ELB should be in a public subnet.

When the user registers a multi-homed instance (an instance that has an Elastic Network Interface (ENI) attached) with a load balancer, the load balancer will route the traffic to the IP address of the primary network interface (eth0).

Reference: <http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/gs-ec2VPC.html>

NEW QUESTION 61

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

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

Answer: A

Explanation:

In AWS OpsWorks, you can control user access to a stack's instances by creating it in a virtual private cloud (VPC). For example, you might not want users to have direct access to your stack's app servers or databases and instead require that all public traffic be channeled through an Elastic Load Balancer.

A VPC consists of one or more subnets, each of which contains one or more instances. Each subnet has an associated routing table that directs outbound traffic based on its destination IP address.

Instances within a VPC can generally communicate with each other, regardless of their subnet. Subnets whose instances can communicate with the Internet are referred to as public subnets. Subnets whose instances can communicate only with other instances in the VPC and cannot communicate directly with the Internet are referred to as private subnets.

AWS OpsWorks requires the VPC to be configured so that every instance in the stack, including instances in private subnets, has access to the following endpoints:

The AWS OpsWorks service, <https://opsworks-instance-service.us-east-1.amazonaws.com> . Amazon S3

The package repositories for Amazon Linux or Ubuntu 12.04 LTS, depending on which operating system you specify.

Your app and custom cookbook repositories. Reference:

<http://docs.aws.amazon.com/opsworks/latest/userguide/workingstacks-vpc.html#workingstacks-vpc-basics>

NEW QUESTION 64

What RAID method is used on the Cloud Block Storage back-end to implement a very high level of reliability and performance?

- A. RAID 1 (Mirror)
- B. RAID 5 (Blocks striped, distributed parity)
- C. RAID 10 (Blocks mirrored and striped)
- D. RAID 2 (Bit level striping)

Answer: C

Explanation:

Cloud Block Storage back-end storage volumes employs the RAID 10 method to provide a very high level of reliability and performance.

Reference: http://www.rackspace.com/knowledge_center/product-faq/cloud-block-storage

NEW QUESTION 66

True or False : "In 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 individual 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 individual 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 individual 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 node

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 individual 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.HowAutoDiscoveryWorks.html>

NEW QUESTION 68

An EC2 instance that performs source/destination checks by default is launched in a private VPC subnet. All security, NACL, and routing definitions are configured as expected. A custom NAT instance is launched.

Which of the following must be done for the custom NAT instance to work?

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

Answer: A

Explanation:

Each EC2 instance performs source/destination checks by default. This means that the instance must be the source or destination of any traffic it sends or receives. However, a NAT instance must be able to send and receive traffic when the source or destination is not itself. Therefore, you must disable source/destination checks on the NAT instance.

Reference:

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_NAT_Instance.html#EIP_Disable_SrcDestCheck

NEW QUESTION 72

How can a user list the IAM Role configured as a part of the launch config?

- A. `as-describe-launch-configs --iam-profile`
- B. `as-describe-launch-configs --show-long`
- C. `as-describe-launch-configs --iam-role`
- D. `as-describe-launch-configs --role`

Answer: B

Explanation:

`as-describe-launch-configs` describes all the launch config parameters created by the AWS account in the specified region. Generally it returns values, such as Launch Config name, Instance Type and AMI ID. If the user wants additional parameters, such as the IAM Profile used in the config, he has to run command: `as-describe-launch-configs --show-long`

NEW QUESTION 73

When using string conditions within IAM, short versions of the available comparators can be used instead of the more verbose ones. `streql` is the short version of the string condition.

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

Answer: A

Explanation:

When using string conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions. For instance, `streql` is the short version of `StringEqualsIgnoreCase` that checks for the exact match between two strings ignoring their case.

Reference: <http://awsdocs.s3.amazonaws.com/SNS/20100331/sns-gsg-2010-03-31.pdf>

NEW QUESTION 74

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

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

Answer: B

Explanation:

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

NEW QUESTION 75

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 79

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

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

Answer: C

Explanation:

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

Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/using_iam.html

NEW QUESTION 82

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 86

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 90

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

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

Answer: A

Explanation:

A service account is a special Google account that can be used by applications to access Google services programmatically. This account belongs to your application or a virtual machine (VM), instead of to an individual end user. Your application uses the service account to call the Google API of a service, so that the users aren't directly involved.

A service account can have zero or more pairs of service account keys, which are used to authenticate to Google. A service account key is a public/private keypair generated by Google. Google retains the public key, while the user is given the private key.

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

NEW QUESTION 95

IAM users do not have permission to create Temporary Security Credentials for federated users and roles by default. In contrast, IAM users can call without the need of any special permissions

- A. GetSessionName
- B. GetFederationToken
- C. GetSessionToken
- D. GetFederationName

Answer: C

Explanation:

Currently the STS API command GetSessionToken is available to every IAM user in your account without previous permission. In contrast, the GetFederationToken command is restricted and explicit permissions need to be granted so a user can issue calls to this particular Action

Reference: <http://docs.aws.amazon.com/STS/latest/UsingSTS/STSPermission.html>

NEW QUESTION 98

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 is automatically a Dedicated instance, regardless of the tenancy of the instance.
- D. None of these are true

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 101

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. Classful 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.html

NEW QUESTION 103

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

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

Answer: A

Explanation:

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

This is helpful when one VPC from the same or different AWS account wants to connect with resources of the other VPC.

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

NEW QUESTION 106

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 110

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 114

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

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

Answer: D

Explanation:

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

NEW QUESTION 119

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

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

Answer: D

Explanation:

The NotAction element lets you specify an exception to a list of actions. Reference:

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

NEW QUESTION 120

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

- A. aws:MultiFactorAuthAge
- B. aws:MultiFactorAuthLast
- C. aws:MFAAge
- D. aws:MultiFactorAuthPrevious

Answer: A

Explanation:

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

NEW QUESTION 122

A user is configuring MySQL RDS with PIOPS. What should be the minimum PIOPS that the user should provision?

- A. 1000
- B. 200
- C. 2000
- D. 500

Answer: A

Explanation:

If a user is trying to enable PIOPS with MySQL RDS, the minimum size of storage should be 100 GB and the minimum PIOPS should be 1000.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.html

NEW QUESTION 125

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

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

Answer: A

Explanation:

Your application does not need to reside in the same VPC as the CloudHSM instance. However, the server or instance on which your application and the HSM client is running must have network (IP) reachability to the HSM. You can establish network connectivity in a variety of ways, including operating your application in the same VPC, with VPC peering, with a VPN connection, or with Direct Connect.
Reference: <https://aws.amazon.com/cloudhsm/faqs/>

NEW QUESTION 126

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

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

Answer: A

Explanation:

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

NEW QUESTION 130

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 135

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

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

Answer: C

Explanation:

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

NEW QUESTION 137

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 142

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

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

Answer: A

Explanation:

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

NEW QUESTION 144

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

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

Answer: C

Explanation:

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

NEW QUESTION 148

Can Provisioned IOPS be used on RDS instances launched in a VPC?

- A. Yes, they can be used only with Oracle based instances.
- B. Yes, they can be used for all RDS instances.
- C. No
- D. Yes, they can be used only with MySQL based instance

Answer: B

Explanation:

The basic building block of Amazon RDS is the DB instance. DB instance storage comes in three types: Magnetic, General Purpose (SSD), and Provisioned IOPS (SSD). When you buy a server, you get CPU, memory, storage, and IOPS, all bundled together. With Amazon RDS, these are split apart so that you can scale them independently. So, for example, if you need more CPU, less IOPS, or more storage, you can easily allocate them.
Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/RDSFAQ.PIOPS.html>

NEW QUESTION 151

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

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

Answer: B

Explanation:

Task Runner polls for tasks and then performs those tasks.

Reference: <http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/what-is-datapipeline.html>

NEW QUESTION 152

A user has created a VPC with public and private subnets using the VPC wizard. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.0.0/24 . The NAT instance ID is i-a12345. Which of the below mentioned entries are required in the main route table attached with the private subnet to allow instances to connect with the internet?

- A. Destination: 20.0.0.0/0 and Target: 80
- B. Destination: 20.0.0.0/0 and Target: i-a12345
- C. Destination: 20.0.0.0/24 and Target: i-a12345
- D. Destination: 0.0.0.0/0 and Target: i-a12345

Answer: D

Explanation:

A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the Internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create two route tables and attach to the subnets. The main route table will have the entry "Destination: 0.0.0.0/0 and Target: i-a12345", which allows all the instances in the private subnet to connect to the internet using NAT.

Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

NEW QUESTION 153

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 154

Which of the following statements is correct about the number of security groups and rules applicable for an EC2-Classic instance and an EC2-VPC network interface?

- A. In EC2-Classic, you can associate an instance with up to 5 security groups and add up to 50 rules to a security group
- B. In EC2-VPC, you can associate a network interface with up to 500 security groups and add up to 100 rules to a security group.
- C. In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 50 rules to a security group
- D. In EC2-VPC, you can associate a network interface with up to 5 security groups and add up to 100 rules to a security group.
- E. In EC2-Classic, you can associate an instance with up to 5 security groups and add up to 100 rules to a security group
- F. In EC2-VPC, you can associate a network interface with up to 500 security groups and add up to 50 rules to a security group.
- G. In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 100 rules to a security group
- H. In EC2-VPC, you can associate a network interface with up to 5 security groups and add up to 50 rules to a security group.

Answer: D

Explanation:

A security group acts as a virtual firewall that controls the traffic for one or more instances. When you launch an instance, you associate one or more security groups with the instance. You add rules to each security group that allow traffic to or from its associated instances. If you're using EC2-Classic, you must use security groups created specifically for EC2-Classic. In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 100 rules to a security group. If you're using EC2-VPC, you must use security groups created specifically for your VPC. In EC2-VPC, you can associate a network interface with up to 5 security groups and add up to 50 rules to a security group.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html>

NEW QUESTION 159

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

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

Answer: A

Explanation:

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

Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html#Sid

NEW QUESTION 160

In Amazon ElastiCache, which of the following statements is correct?

- A. When you launch an ElastiCache cluster into an Amazon VPC private subnet, every cache node is assigned a public IP address within that subnet.
- B. You cannot use ElastiCache in a VPC that is configured for dedicated instance tenancy.
- C. If your AWS account supports only the EC2-VPC platform, ElastiCache will never launch your cluster in a VPC.
- D. ElastiCache is not fully integrated with Amazon Virtual Private Cloud (VPC).

Answer: B

Explanation:

The VPC must allow non-dedicated EC2 instances. You cannot use ElastiCache in a VPC that is configured for dedicated instance tenancy.

Reference: <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/AmazonVPC.EC.html>

NEW QUESTION 165

An organization has setup RDS with VPC. The organization wants RDS to be accessible from the internet. Which of the below mentioned configurations is not required in this scenario?

- A. The organization must enable the parameter in the console which makes the RDS instance publicly accessible.
- B. The organization must allow access from the internet in the RDS VPC security group,
- C. The organization must setup RDS with the subnet group which has an external IP.
- D. The organization must enable the VPC attributes DNS hostnames and DNS resolution

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, 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 which the user assigns to the RDS DB instances. A DB subnet group allows the user to specify a particular VPC when creating DB instances. If the RDS instance is required to be accessible from the internet:

The organization must setup that the RDS instance is enabled with the VPC attributes, DNS hostnames and DNS resolution.

The organization must enable the parameter in the console which makes the RDS instance publicly accessible.

The organization must allow access from the internet in the RDS VPC security group. Reference:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.html

NEW QUESTION 170

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

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

Answer: C

NEW QUESTION 171

To serve Web traffic for a popular product your chief financial officer and IT director have purchased 10 ml large heavy utilization Reserved Instances (RIs) evenly spread across two availability zones: Route 53 is used to deliver the traffic to an Elastic Load Balancer (ELB). After several months, the product grows even more popular and you need additional capacity As a result, your company purchases two C3.2xlarge medium utilization Ris You register the two c3 2xlarge instances with your ELB and quickly find that the ml large instances are at 100% of capacity and the c3 2xlarge instances have significant capacity that's unused Which option is the most cost effective and uses EC2 capacity most effectively?

- A. Configure Autoscaling group and Launch Configuration with ELB to add up to 10 more on-demand m1 .large instances when triggered by Cloudwatc
- B. Shut off c3.2xlarge instances.
- C. Configure ELB with two c3.2xlarge instances and use on-demand Autoscaling group for up to two additional c3.2xlarge instance
- D. Shut off m1 .large instances.
- E. Route traffic to EC2 m1 .large and c3.2xlarge instances directly using Route 53 latency based routing and health check
- F. Shut off ELB.
- G. Use a separate ELB for each instance type and distribute load to ELBs with Route 53 weighted round robin.

Answer: B

NEW QUESTION 176

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

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

Answer: A

NEW QUESTION 179

You are designing a photo-sharing mobile app. The application will store all pictures in a single Amazon S3 bucket. Users will upload pictures from their mobile device directly to Amazon S3 and will be able to view and download their own pictures directly from Amazon S3. You want to configure security to handle potentially millions of users in the most secure manner possible. What should your server-side application do when a new user registers on the photo-sharing mobile application?

- A. Create an IAM use
- B. Update the bucket policy with appropriate permissions for the IAM use
- C. Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.
- D. Create an IAM use
- E. Assign appropriate permissions to the IAM use
- F. Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.
- G. Create a set of long-term credentials using AWS Security Token Service with appropriate permission
- H. Store these credentials in the mobile app and use them to access Amazon S3.
- I. Record the user's information in Amazon RDS and create a role in IAM with appropriate permission
- J. When the user uses their mobile app, create temporary credentials using the AWS Security Token Service "AssumeRole" functio
- K. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.
- L. Record the user's information in Amazon DynamoD
- M. When the user uses their mobile app, create temporary credentials using AWS Security Token Service with appropriate permission
- N. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.

Answer: D

NEW QUESTION 184

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 c3.8xlarge 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 3x1TB 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 189

A large real-estate brokerage is exploring the option o(adding a cost-effective location based alert to their existing mobile application The application backend infrastructure currently runs on AWS Users who opt in to this service will receive alerts on their mobile device regarding real-estate otters in proximity to their location. For the alerts to be relevant delivery time needs to be in the low minute count the existing mobile app has 5 million users across the US. Which one of the following architectural suggestions would you make to the customer?

- A. The mobile application will submit its location to a web service endpoint utilizing Elastic Load Balancing and EC2 instances: DynamoDB will be used to store and retrieve relevant offers EC2 instances will communicate with mobile earners/device providers to push alerts back to mobile application.
- B. Use AWS DirectConnect or VPN to establish connectMty with mobile carriers EC2 instances will receive the mobile applications ' location through carrier connection: RDS will be used to store and relevant offers EC2 instances will communicate with mobile carriers to push alerts back to the mobile application
- C. The mobile application will send device location using SQ
- D. EC2 instances will retrieve the relevant others from DynamoDB AWS Mobile Push will be used to send offers to the mobile application
- E. The mobile application will send device location using AWS Nmobile Push EC2 instances will retrieve the relevant offers from DynamoDB EC2 instances will communicate with mobile carriers/device providers to push alerts back to the mobile application.

Answer: A

NEW QUESTION 194

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

Answer: A

NEW QUESTION 198

An AWS customer is deploying an application that is composed of an AutoScaling group of EC2 Instances.

The customer's security policy requires that every outbound connection from these instances to any other service within the customer's Virtual Private Cloud must be authenticated using a unique x.509 certificate that contains the specific instance-id.

In addition, x.509 certificates must be 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 the 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 Service.
- 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 (that contains the specific instance-id).

Answer: A

NEW QUESTION 200

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-tier architecture and currently uses MySQL 5.6 for data persistence
- 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 the HQ region
- D. For each regional deployment, use MySQL 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 203

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

What AWS architecture would you recommend?

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

Answer: A

NEW QUESTION 204

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

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

Answer: AC

NEW QUESTION 207

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 212

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

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

Answer: C

NEW QUESTION 213

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 VMDK is almost full; The virtual network interface still uses the 10Gbps driver, which leaves your 100Mbps WAN connection completely underutilized;

It is currently running on a highly customized Windows VM within a VMware environment; You do not have the installation media;

This is a mission critical application with an RTO (Recovery Time Objective) of 8 hours. RPO (Recovery Point Objective) of 1 hour. How could you best migrate this application to AWS while meeting your business continuity requirements?

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

Answer: A

NEW QUESTION 215

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

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

Answer: C

NEW QUESTION 217

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 of 10 minutes where all 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 and its 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 221

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

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

Answer: B

NEW QUESTION 224

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

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

Which design would you choose to meet these requirements?

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

Answer: A

NEW QUESTION 228

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

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

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

Answer: C

NEW QUESTION 233

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

Which architecture change could you introduce to reduce costs and still keep your web application secure and scalable?

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

Answer: C

NEW QUESTION 235

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

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

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

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

Answer: B

NEW QUESTION 240

You are designing an S3 solution that requires HTTPS clients to be authenticated by the Webserver using client certificate authentication. The solution must be resilient.

Which of the following options would you consider for configuring the web server infrastructure? (Choose 2 answers)

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

Answer: AB

NEW QUESTION 244

You are designing a connectivity solution between on-premises infrastructure and Amazon VPC. Your servers on-premises will be communicating with your VPC instances. You will be establishing IPsec tunnels over the Internet. You will be using VPN gateways, and terminating the IPsec tunnels on AWS supported customer gateways.

Which of the following objectives would you achieve by implementing an IPsec tunnel as outlined above? Choose 4 answers

- A. End-to-end protection of data in transit
- B. End-to-end Identity authentication
- C. Data encryption across the Internet
- D. Protection of data in transit over the Internet
- E. Peer identity authentication between VPN gateway and customer gateway
- F. Data integrity protection across the Internet

Answer: CDEF

NEW QUESTION 245

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

What should you do in order to avoid errors for future deployments? (Choose 2 answers)

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

Answer: CD

NEW QUESTION 247

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

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

Answer: B

NEW QUESTION 252

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

```
./marketing
./sales
.Hegal
```

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

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

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

- A. True
- B. False

Answer: B

NEW QUESTION 256

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

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

Answer: A

NEW QUESTION 257

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

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

Answer: B

NEW QUESTION 262

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

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

Answer: C

NEW QUESTION 266

You've been brought in as solutions architect to assist an enterprise customer with their migration of an e-commerce platform to Amazon Virtual Private Cloud (VPC) The previous architect has already deployed a 3-tier VPC.

The configuration is as follows: VPC: vpc-2f8bc447

IGW: igw-2d8bc445 NACL: ad-208bc448

Subnets and Route Tables: Web servers: subnet-258bc44d

Application servers: subnet-248bc44c Database servers: subnet-9189c6f9 Route Tables:

rtb-218bc449 rtb-238bc44b Associations:

subnet-258bc44d : rtb-218bc449 subnet-248bc44c : rtb-238bc44b subnet-9189c6f9 : rtb-238bc44b

You are now ready to begin deploying EC2 instances into the VPC Web servers must have direct access to the internet Application and database servers cannot have direct access to the internet.

Which configuration below will allow you the ability to remotely administer your application and database servers, as well as allow these servers to retrieve updates from the Internet?

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

Answer: A

NEW QUESTION 267

You are designing a multi-platform web application for AWS The application will run on EC2 instances and will be accessed from PCs, tablets and smart phones Supported accessing platforms are Windows, MacOS, IOS and Android Separate sticky session and SSL certificate setups are required for different platform types which of the following describes the most cost effective and performance efficient architecture setup?

- A. Setup a hybrid architecture to handle session state and SSL certificates on-prem and separate EC2 Instance groups running web applications for different platform types running in a VPC.
- B. Set up one ELB for all platforms to distribute load among multiple instance under it Each EC2 instance implements all functionality for a particular platform.
- C. Set up two ELBs The first ELB handles SSL certificates for all platforms and the second ELB handles session stickiness for all platforms for each ELB run separate EC2 instance groups to handle the web application for each platform.
- D. Assign multiple ELBs to an EC2 instance or group of EC2 instances running the common components of the web application, one ELB for each platform type Session stickiness and SSL termination are done at the ELBs.

Answer: D

NEW QUESTION 268

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

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

Answer: D

NEW QUESTION 269

A benefits enrollment company is hosting a 3-tier web application running in a VPC on AWS which includes a NAT (Network Address Translation) instance in the public Web tier. There is enough provisioned capacity for the expected workload for the new fiscal year benefit enrollment period plus some extra overhead Enrollment proceeds nicely for two days and then the web tier becomes unresponsive, upon investigation using CloudWatch and other monitoring tools it is discovered that there is an extremely large and unanticipated amount of inbound traffic coming from a set of 15 specific IP addresses over port 80 from a country where the benefits company has no customers. The web tier instances are so overloaded that benefit enrollment administrators cannot even SSH into them. Which actMty would be useful in defending against this attack?

- A. Create a custom route table associated with the web tier and block the attacking IP addresses from the IGW (Internet Gateway)
- B. Change the EIP (Elastic IP Address) of the NAT instance in the web tier subnet and update the Nlain Route Table with the new EIP
- C. Create 15 Security Group rules to block the attacking IP addresses over port 80
- D. Create an inbound NACL (Network Access control list) associated with the web tier subnet with deny rules to block the attacking IP addresses

Answer: D

NEW QUESTION 274

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 276

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

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

Answer: C

NEW QUESTION 281

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

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

Answer: A

NEW QUESTION 283

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

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

Answer: C

NEW QUESTION 286

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

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

Answer: AE

NEW QUESTION 291

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

- A. Implement third party volume encryption tools
- B. Implement SSL/TLS for all services running on the server
- C. Encrypt data inside your applications before storing it on EBS
- D. Encrypt data using native data encryption drivers at the file system level
- E. Do nothing as EBS volumes are encrypted by default

Answer: ACD

NEW QUESTION 294

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

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

Answer: D

NEW QUESTION 296

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 300

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

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

Answer: D

NEW QUESTION 304

You are designing Internet connectivity for your VPC. The Web servers must be available on the Internet. The application must have a highly available architecture. Which alternatives should you consider? (Choose 2 answers)

- A. Configure a NAT instance in your VPC. Create a default route via the NAT instance and associate it with all subnets. Configure a DNS A record that points to the NAT instance public IP address.
- B. Configure a CloudFront distribution and configure the origin to point to the private IP addresses of your Web servers. Configure a Route53 CNAME record to your CloudFront distribution.
- C. Place all your web servers behind ELB. Configure a Route53 CNAME to point to the ELB DNS name.
- D. Assign EIPs to all web servers
- E. Configure a Route53 record set with all EIPs, with health checks and DNS failover.
- F. Configure ELB with an EIP. Place all your Web servers behind ELB. Configure a Route53 A record that points to the EIP.

Answer: CD

NEW QUESTION 306

Your company has an on-premises multi-tier PHP web application, which recently experienced downtime due to a large burst in web traffic due to a company announcement. Over the coming days, you are expecting similar announcements to drive similar unpredictable bursts, and are looking to find ways to quickly improve your infrastructure's ability to handle unexpected increases in traffic.

The application currently consists of 2 tiers: a web tier which consists of a load balancer and several Linux Apache web servers as well as a database tier which hosts a Linux server hosting a MySQL database. Which scenario below will provide full site functionality, while helping to improve the ability of your application in the short timeframe required?

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

Answer: C

NEW QUESTION 310

An ERP application is deployed across multiple AZs in a single region. In the event of failure, the Recovery Time Objective (RTO) must be less than 3 hours, and the Recovery Point Objective (RPO) must be 15 minutes. The customer realizes that data corruption occurred roughly 1.5 hours ago. What DR strategy could be used to achieve this RTO and RPO in the event of this kind of failure?

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

Answer: A

NEW QUESTION 312

You control access to S3 buckets and objects with:

- A. Identity and Access Management (IAM) Policies.
- B. Access Control Lists (ACLs).

- C. Bucket Policies.
- D. All of the above

Answer: D

NEW QUESTION 313

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 315

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

What is the problem and a valid solution?

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

Answer: C

NEW QUESTION 320

.....

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>