

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

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

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

Answer: B

Explanation:

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

Reference:

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

NEW QUESTION 3

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

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

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 4

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 5

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

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

Answer: D

Explanation:

To submit an AWS Direct Connect connection request, you need to provide the following information: Your contact information.
The AWS Direct Connect Location to connect to.

Details of AWS Direct Connect partner if you use the AWS Partner Network (APN) service. The port speed you require, either 1 Gbps or 10 Gbps.
Reference: <http://docs.aws.amazon.com/directconnect/latest/UserGuide/getstarted.html#ConnectionRequest>

NEW QUESTION 6

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 7

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

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

Answer: C

Explanation:

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

NEW QUESTION 8

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

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

Answer: C

Explanation:

RDS DB instance storage comes in two types: standard and provisioned IOPS. Standard storage is allocated on the Amazon EBS volumes and connected to the user's DB instance. Provisioned IOPS uses

optimized EBS volumes and an optimized configuration stack. It provides additional, dedicated capacity for the EBS I/O.

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

NEW QUESTION 9

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

- A. ArnCheck

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

Answer: D

Explanation:

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

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

NEW QUESTION 10

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

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 10

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 12

IV|apMySite is setting up a web application in the AWS VPC. The organization has decided to use an AWS RDS instead of using its own DB instance for HA and DR requirements.

The organization also wants to secure RDS access. How should the web application be setup with RDS?

- A. Create a VPC with one public and one private subnet
- B. Launch an application instance in the public subnet while RDS is launched in the private subnet.
- C. Setup a public and two private subnets in different AZs within a VPC and create a subnet group
- D. Launch RDS with that subnet group.
- E. Create a network interface and attach two subnets to it
- F. Attach that network interface with RDS while launching a DB instance.
- G. Create two separate VPCs and launch a Web app in one VPC and RDS in a separate VPC and connect them with VPC peering.

Answer: B

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources, 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 the security and operational needs.

A DB subnet group is a collection of subnets (generally private) that a user can create in a VPC and assign to the RDS DB instances. A DB subnet group allows the user to specify a particular VPC when creating the DB instances. Each DB subnet group should have subnets in at least two Availability Zones in a given region.

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

NEW QUESTION 14

The Statement element, of an AWS IAM policy, contains an array of individual statements. Each individual statement is a(n) block enclosed in braces { }.

- A. XML
- B. JavaScript
- C. JSON
- D. AJAX

Answer: C

Explanation:

The Statement element, of an IAM policy, contains an array of individual statements. Each individual statement is a JSON block enclosed in braces { }.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html

NEW QUESTION 15

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

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

Answer: A

Explanation:

Multiple compute resources can be used on the same pipeline in AWS Data Pipeline by defining multiple cluster objects in your definition file and associating the cluster to use for each actMty via its runsOn field, which allows pipelines to combine AWS and on-premise resources, or to use a mix of instance types for their actMtys.

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

NEW QUESTION 19

One of your AWS Data Pipeline actMtys 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 21

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 25

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 dat
- C. Thus, whenever the instance is booted the script assigns the fixed MAC address to that instance.
- D. The instance MAC address never change
- E. Thus, it is not required to register the MAC address every time.
- F. AWS never provides a MAC address to an instance; instead the instance ID is used for identifying the instance for any software registration.

Answer: A

Explanation:

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

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

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

NEW QUESTION 31

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 34

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

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

Answer: C

Explanation:

An AWS Direct Connect location provides access to Amazon Web Services in the region it is associated with, as well as access to other US regions. For example, you can provision a single connection to any AWS Direct Connect location in the US and use it to access public AWS services in all US Regions and AWS GovCloud (US).

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

NEW QUESTION 36

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 39

What types of identities do Amazon Cognito identity pools support?

- A. They support both authenticated and unauthenticated identities.
- B. They support only unauthenticated identities.
- C. They support neither authenticated nor unauthenticated identities.
- D. They support only authenticated identities.

Answer: A

Explanation:

Amazon Cognito identity pools support both authenticated and unauthenticated identities. Authenticated identities belong to users who are authenticated by a public login provider or your own backend authentication process. Unauthenticated identities typically belong to guest users. Reference: <http://docs.aws.amazon.com/cognito/devguide/identity/identity-pool/>

NEW QUESTION 40

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

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

Answer: A

Explanation:

Temporary credentials in IAM are valid throughout their defined duration of time and hence can't be revoked. However, because permissions are evaluated each time an AWS request is made using the credentials, you can achieve the effect of revoking the credentials by changing the permissions for the credentials even after they have been issued. Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp_control-access_disable-perms.html

NEW QUESTION 44

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

- A. "Effect": "Allow", "Action": ["Describe"], "Resource": "Billing"
- B. "Effect": "Allow", "Action": ["aws-portal:ViewBilling"], "Resource": "**"
- C. "Effect": "Allow", "Action": ["aws-portal:ViewUsage"], "Resource": "**"
- D. "Effect": "Allow", "Action": ["AccountUsage"], "Resource": "**"

Answer: C

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the CFO wants to allow only AWS usage report page access, the policy for that IAM user will be as given below:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow", "Action": [
        "aws-portal:ViewUsage"
      ],
      "Resource": "**"
    }
  ]
}
```

Reference: <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/billing-permissions-ref.html>

NEW QUESTION 46

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

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

Answer: B

Explanation:

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

NEW QUESTION 49

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

How can the organization achieve this by running web server on a single instance?

- A. It is not possible to have two IP addresses for a single instance.
- B. The organization should create two network interfaces with the same subnet and security group to assign separate IPs to each network interface.
- C. The organization should create two network interfaces with separate subnets so one instance can have two subnets and the respective security groups for

controlled access.

D. The organization should launch an instance with two separate subnets using the same network interface which allows to have a separate CIDR as well as security groups.

Answer: C

Explanation:

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

The user can create a management network using two separate network interfaces. For the present scenario it is required that the secondary network interface on the instance handles the public facing traffic and the primary network interface handles the back-end management traffic and it is connected to a separate subnet in the VPC that has more restrictive access controls. The public facing interface, which may or may not be behind a load balancer, has an associated security group to allow access to the server from the internet while the private facing interface has an associated security group allowing SSH access only from an allowed range of IP addresses either within the VPC or from the internet, a private subnet within the VPC or a virtual private gateway.

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

NEW QUESTION 53

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 54

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

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

Answer: C

Explanation:

A Provisioned IOPS (SSD) volume can range in size from 4 GiB to 16 TiB and you can provision up to 20,000 IOPS per volume.

Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html#EBSVolumeTypes_piops

NEW QUESTION 55

Which of the following cache engines does Amazon ElastiCache support?

- A. Amazon ElastiCache supports Memcached and Redis.
- B. Amazon ElastiCache supports Redis and WinCache.
- C. Amazon ElastiCache supports Memcached and Hazelcast.
- D. Amazon ElastiCache supports Memcached only

Answer: A

Explanation:

The cache engines supported by Amazon ElastiCache are Memcached and Redis.

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

NEW QUESTION 57

In a VPC, can you modify a set of DHCP options after you create them?

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

Answer: C

Explanation:

After you create a set of DHCP options, you can't modify them. If you want your VPC to use a different set of DHCP options, you must create a new set and associate them with your VPC. You can also set up your VPC to use no DHCP options at all.

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

NEW QUESTION 62

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

A. 24

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

Answer: B

Explanation:

By default, temporary security credentials for an IAM user are valid for a maximum of 12 hours, but you can request a duration as short as 15 minutes or as long as 36 hours.

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

NEW QUESTION 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 68

An organization is setting up a highly scalable application using Elastic Beanstalk. They are using Elastic Load Balancing (ELB) as well as a Virtual Private Cloud (VPC) with public and private subnets. They have the following requirements:

- . All the EC2 instances should have a private IP
- . All the EC2 instances should receive data via the ELB's. Which of these will not be needed in this setup?

- A. Launch the EC2 instances with only the public subnet.
- B. Create routing rules which will route all inbound traffic from ELB to the EC2 instances.
- C. Configure ELB and NAT as a part of the public subnet only.
- D. Create routing rules which will route all outbound traffic from the EC2 instances through NA

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. If the organization wants the Amazon EC2 instances to have a private IP address, he should create a public and private subnet for VPC in each Availability Zone (this is an AWS Elastic Beanstalk requirement). The organization should add their public resources, such as ELB and NAT to the public subnet, and AWS Elastic Beanstalk will assign them unique elastic IP addresses (a static, public IP address). The organization should launch Amazon EC2 instances in a private subnet so that AWS Elastic Beanstalk assigns them non-routable private IP addresses. Now the organization should configure route tables with the following rules:

- . route all inbound traffic from ELB to EC2 instances
- . route all outbound traffic from EC2 instances through NAT

Reference: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/AWSHowTo-vpc.html>

NEW QUESTION 73

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 77

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

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

Answer: B

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. With VPC the user can specify multiple private IP addresses for his instances.

The number of network interfaces and private IP addresses that a user can specify for an instance depends on the instance type. With each network interface the organization can assign an EIP. This scenario helps when the user wants to host multiple websites on a single EC2 instance by using multiple SSL certificates on a single server and associating each certificate with a specific EIP address. It also helps in scenarios for operating network appliances, such as firewalls or load balancers that have multiple private IP addresses for each network interface.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/NIultipleIP.html>

NEW QUESTION 82

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 87

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

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

Answer: C

Explanation:

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

Instead they should have two separate DBs in each site and setup data replication based on RTO(recovery time objective) of the organization.

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

NEW QUESTION 90

Select the correct statement about Amazon ElastiCache.

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

Answer: A

Explanation:

ElastiCache is a web service that makes it easy to set up, manage, and scale a distributed in-memory cache environment in the cloud. It provides a high-performance, scalable, and cost-effective caching solution, while removing the complexity associated with deploying and managing a distributed cache environment. With ElastiCache, you can quickly deploy your cache environment, without having to provision hardware or install software.

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

NEW QUESTION 92

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 94

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 99

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 100

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

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

Answer: A

Explanation:

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

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

Reference:

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

NEW QUESTION 104

True or False: The Amazon ElastiCache clusters are not available for use in VPC at this time.

- A. TRUE
- B. True, but they are available only in the GovCloud.
- C. True, but they are available only on request.
- D. FALSE

Answer: D

Explanation:

Amazon ElastiCache clusters can be run in an Amazon VPC. With Amazon VPC, you can define a virtual network topology and customize the network configuration to closely resemble a traditional network that you might operate in your own datacenter. You can now take advantage of the manageability, availability and scalability benefits of Amazon ElastiCache Clusters in your own isolated network. The same functionality of Amazon ElastiCache, including automatic failure detection, recovery, scaling, auto discovery, Amazon CloudWatch metrics, and software patching, are now available in Amazon VPC. Reference: <http://aws.amazon.com/about-aws/whats-new/2012/12/20/amazon-elasticache-announces-support-for-a-mazon-vpc/>

NEW QUESTION 106

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 111

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 112

You are setting up some EBS volumes for a customer who has requested a setup which includes a RAID (redundant array of inexpensive disks). AWS has some recommendations for RAID setups. Which RAID setup is not recommended for Amazon EBS?

- A. RAID 1 only
- B. RAID 5 only
- C. RAID 5 and RAID 6
- D. RAID 0 only

Answer: C

Explanation:

With Amazon EBS, you can use any of the standard RAID configurations that you can use with a traditional bare metal server, as long as that particular RAID configuration is supported by the operating system for your instance. This is because all RAID is accomplished at the software level. For greater I/O performance than you can achieve with a single volume, RAID 0 can stripe multiple volumes together; for on-instance redundancy, RAID 1 can mirror two volumes together. RAID 5 and RAID 6 are not recommended for Amazon EBS because the parity write operations of these RAID modes consume some of the IOPS available to your volumes.

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

NEW QUESTION 116

Who is responsible for modifying the routing tables and networking ACLs in a VPC to ensure that a DB instance is reachable from other instances in the VPC?

- A. AWS administrators
- B. The owner of the AWS account
- C. Amazon
- D. The DB engine vendor

Answer: B

Explanation:

You are in charge of configuring the routing tables of your VPC as well as the network ACLs rules needed to make your DB instances accessible from all the instances of your VPC that need to communicate with it.

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

NEW QUESTION 119

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_pio ps

NEW QUESTION 123

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

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

Answer: D

Explanation:

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

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

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html>

NEW QUESTION 126

A government client needs you to set up secure cryptographic key storage for some of their extremely confidential data. You decide that the AWS CloudHSM is the best service for this. However, there seem to be a few pre-requisites before this can happen, one of those being a security group that has certain ports open. Which of the following is correct in regards to those security groups?

- A. A security group that has no ports open to your network.
- B. A security group that has only port 3389 (for RDP) open to your network.
- C. A security group that has only port 22 (for SSH) open to your network.
- D. A security group that has port 22 (for SSH) or port 3389 (for RDP) open to your network

Answer: D

Explanation:

AWS CloudHSM provides secure cryptographic key storage to customers by making hardware security modules (HSMs) available in the AWS cloud.

AWS CloudHSM requires the following environment before an HSM appliance can be provisioned. A virtual private cloud (VPC) in the region where you want the AWS CloudHSM service.

One private subnet (a subnet with no Internet gateway) in the VPC. The HSM appliance is provisioned into this subnet.

One public subnet (a subnet with an Internet gateway attached). The control instances are attached to this subnet.

An AWS Identity and Access Management (IAM) role that delegates access to your AWS resources to AWS CloudHSM.

An EC2 instance, in the same VPC as the HSM appliance, that has the SafeNet client software installed. This instance is referred to as the control instance and is used to connect to and manage the HSM appliance.

A security group that has port 22 (for SSH) or port 3389 (for RDP) open to your network. This security group is attached to your control instances so you can access them remotely.

NEW QUESTION 130

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 132

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

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

Answer: A

Explanation:

You can initialize Amazon Mobile Analytics using AWS IAM accounts. AWS recommends using Amazon Cognito for security best practices.

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

NEW QUESTION 133

Mike is appointed as Cloud Consultant in ExamKiller.com. ExamKiller has the following VPCs set-up in the US East Region:

A VPC with CIDR block 10.10.0.0/16, a subnet in that VPC with CIDR block 10.10.1.0/24 A VPC with CIDR block 10.40.0.0/16, a subnet in that VPC with CIDR block 10.40.1.0/24

ExamKiller.com is trying to establish network connection between two subnets, a subnet with CIDR block 10.10.1.0/24 and another subnet with CIDR block 10.40.1.0/24. Which one of the following solutions should Mike recommend to ExamKiller.com?

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

Answer: C

Explanation:

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

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

NEW QUESTION 134

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 136

An organization is setting up their website on AWS. The organization is working on various security measures to be performed on the AWS EC2 instances. Which of the below mentioned security mechanisms will not help the organization to avoid future data leaks and identify security weaknesses?

- A. Run penetration testing on AWS with prior approval from Amazon.
- B. Perform SQL injection for application testing.
- C. Perform a Code Check for any memory leaks.
- D. Perform a hardening test on the AWS instanc

Answer: C

Explanation:

AWS security follows the shared security model where the user is as much responsible as Amazon. Since Amazon is a public cloud it is bound to be targeted by hackers. If an organization is planning to host their application on AWS EC2, they should perform the below mentioned security checks as a measure to find any security weakness/data leaks:

Perform penetration testing as performed by attackers to find any vulnerability. The organization must take an approval from AWS before performing penetration testing

Perform hardening testing to find if there are any unnecessary ports open Perform SQL injection to find any DB security issues

The code memory checks are generally useful when the organization wants to improve the application performance.

Reference: <http://aws.amazon.com/security/penetration-testing/>

NEW QUESTION 138

In Amazon ElastiCache, the default cache port is:

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

Answer: D

Explanation:

In Amazon ElastiCache, you can specify a new port number for your cache cluster, which by default is 11211 for Memcached and 6379 for Redis.

Reference: <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/GettingStarted.AuthorizeAccess.htm>

NEW QUESTION 143

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 145

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 150

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

- A. Use OAuth 2.0 to retrieve temporary AWS security credentials to enable your NOC members to sign in to the AWS Management Console.
- B. Use web Identity Federation to retrieve AWS temporary security credentials to enable your NOC members to sign in to the AWS 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 SAML 2.0-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 153

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 Billing
- F. This will give IAM users in the Master account access to resources in the Dev and Test accounts

Answer: C

NEW QUESTION 155

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 Cloudwatch
- 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 158

You have deployed a web application targeting a global audience across multiple AWS Regions under the domain name.example.com. You decide to use Route53 Latency-Based Routing to serve web requests to users from the region closest to the user. To provide business continuity in the event of server downtime you configure weighted record sets associated with two web servers in separate Availability Zones per region. Dunning a DR test you notice that when you disable all web servers in one of the regions Route53 does not automatically direct all users to the other region. What could be happening? (Choose 2 answers)

- A. Latency resource record sets cannot be used in combination with weighted resource record sets.
- B. You did not setup an HTTP health check to one or more of the weighted resource record sets associated with me disabled web servers.
- C. The value of the weight associated with the latency alias resource record set in the region with the disabled servers is higher than the weight for the other

- region.
- D. One of the two working web servers in the other region did not pass its HTTP health check.
- E. You did not set "Evaluate Target Health" to "Yes" on the latency alias resource record set associated with example.com in the region where you disabled the servers.

Answer: BE

NEW QUESTION 160

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

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

Answer: C

NEW QUESTION 165

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 169

You are the new IT architect in a company that operates a mobile sleep tracking application. When activated at night, the mobile app is sending collected data points of 1 kilobyte every 5 minutes to your backend. The backend takes care of authenticating the user and writing the data points into an Amazon DynamoDB table. Every morning, you scan the table to extract and aggregate last night's data on a per user basis, and store the results in Amazon S3. Users are notified via Amazon SNS mobile push notifications that new data is available, which is parsed and visualized by the mobile app. Currently you have around 100k users who are mostly based out of North America. You have been tasked to optimize the architecture of the backend system to lower cost. What would you recommend? Choose 2 answers

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

Answer: AD

NEW QUESTION 170

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

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

Answer: D

NEW QUESTION 172

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 177

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

- A. You are running the proxy on an undersized EC2 instance type so network throughput is not sufficient for all instances to download their updates in time.
- B. You are running the proxy on a sufficiently-sized EC2 instance in a private subnet 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 179

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

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

Answer: B

NEW QUESTION 182

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

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

Answer: AC

NEW QUESTION 187

You are designing an intrusion detection prevention (IDS/IPS) solution for a customer web application in a single VPC. You are considering the options for implementing IOS IPS protection for traffic coming from the Internet. Which of the following options would you consider? (Choose 2 answers)

- A. Implement IDS/IPS agents on each Instance running In VPC
- B. Configure an instance in each subnet to switch its network interface card to promiscuous mode and analyze network traffic.
- C. Implement Elastic Load Balancing with SSL listeners In front of the web applications
- D. Implement a reverse proxy layer in front of web servers and configure IDS/IPS agents on each reverse proxy server.

Answer: BD

NEW QUESTION 189

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

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

Answer: CEF

NEW QUESTION 194

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

How should they architect their solution?

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

Answer: D

NEW QUESTION 196

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

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

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

Answer: C

NEW QUESTION 201

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

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

Answer: D

NEW QUESTION 204

You have deployed a three-tier web application in a VPC with a CIDR block of 10.0.0.0/28. You initially deploy two web servers, two application servers, two database servers and one NAT instance for a total of seven EC2 instances. The web, application and database servers are deployed across two availability zones (AZs). You also deploy an ELB in front of the two web servers, and use Route53 for DNS. Web traffic gradually increases in the first few days following the deployment, so you attempt to double the number of instances in each tier of the application to handle the new load. Unfortunately, some of these new instances fail to launch.

Which of the following could be the root cause? (Choose 2 answers)

- A. AWS reserves the first and the last private IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances
- B. The Internet Gateway (IGW) of your VPC has scaled-up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches
- C. The ELB has scaled-up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches
- D. AWS reserves one IP address in each subnet's CIDR block for Route53 so you do not have enough addresses left to launch all of the new EC2 instances

E. AWS reserves the first four and the last IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances

Answer: CE

NEW QUESTION 207

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 212

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

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

Answer: D

NEW QUESTION 217

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 220

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 222

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 domain. 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 227

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

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

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

Answer: ABD

NEW QUESTION 232

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 oaten 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 me 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 235

You are designing a data leak prevention solution for your VPC environment. You want your VPC Instances to be able to access software depots and distributions on the Internet for product updates. The depots and distributions are accessible via third party CDNs by their URLs. You want to explicitly deny any other outbound connections from your VPC instances to hosts on the internet.

Which of the following options would you consider?

- A. Configure a web proxy server in your VPC and enforce URL-based rules for outbound access. Remove default routes.
- B. Implement security groups and configure outbound rules to only permit traffic to software depots.
- C. Move all your instances into private VPC subnets remove default routes from all routing tables and add specific routes to the software depots and distributions only.
- D. Implement network access control lists to all specific destinations, with an Implicit deny as a rule.

Answer: A

NEW QUESTION 239

Your company hosts a social media site supporting users in multiple countries. You have been asked to provide a highly available design for the application that leverages multiple regions for the most recently accessed content and latency sensitive portions of the website. The most latency sensitive component of the

application involves reading user preferences to support web site personalization and ad selection. In addition to running your application in multiple regions, which option will support this application's requirements?

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

Answer: A

NEW QUESTION 240

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 245

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

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

Answer: C

NEW QUESTION 250

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 251

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

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

Answer: C

NEW QUESTION 255

You are developing a new mobile application and are considering storing user preferences in AWS. This would provide a more uniform cross-device experience

to users using multiple mobile devices to access the application. The preference data for each user is estimated to be 50KB in size. Additionally, 5 million customers are expected to use the application on a regular basis. The solution needs to be cost-effective, highly available, scalable, and secure. How would you design a solution to meet the above requirements?

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

Answer: B

NEW QUESTION 257

You deployed your company website using Elastic Beanstalk and you enabled log file rotation to S3. An Elastic Map Reduce job is periodically analyzing the logs on S3 to build a usage dashboard that you share with your CIO.

You recently improved overall performance of the website using Cloud Front for dynamic content delivery and your website as the origin.

After this architectural change, the usage dashboard shows that the traffic on your website dropped by an order of magnitude. How do you fix your usage dashboard?

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

Answer: D

NEW QUESTION 261

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 264

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 266

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 267

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 268

You have recently joined a startup company building sensors to measure street noise and air quality in urban areas. The company has been running a pilot deployment of around 100 sensors for 3 months each sensor uploads 1KB of sensor data every minute to a backend hosted on AWS.

During the pilot, you measured a peak of 10 IOPS on the database, and you stored an average of 3GB of sensor data per month in the database.

The current deployment consists of a load-balanced auto scaled Ingestion layer using EC2 instances and a PostgreSQL RDS database with 500GB standard storage.

The pilot is considered a success and your CEO has managed to get the attention of some potential investors. The business plan requires a deployment of at least 100K sensors which needs to be supported by the backend. You also need to store sensor data for at least two years to be able to compare year over year improvements.

To secure funding, you have to make sure that the platform meets these requirements and leaves room for further scaling. Which setup will meet the requirements?

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

Answer: C

NEW QUESTION 272

Your company is in the process of developing a next generation pet collar that collects biometric information to assist families with promoting healthy lifestyles for their pets. Each collar will push 30kb of biometric data in JSON format every 2 seconds to a collection platform that will process and analyze the data providing health trending information back to the pet owners and veterinarians via a web portal. Management has tasked you to architect the collection platform ensuring the following requirements are met.

Provide the ability for real-time analytics of the inbound biometric data. Ensure processing of the biometric data is highly durable, elastic and parallel. The results of the analytic processing should be persisted for data mining.

Which architecture outlined below will meet the initial requirements for the collection platform?

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

Answer: B

NEW QUESTION 277

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 278

You are implementing AWS Direct Connect. You intend to use AWS public service endpoints such as Amazon S3, across the AWS Direct Connect link. You want other Internet traffic to use your existing link to an Internet Service Provider.

What is the correct way to configure AWS Direct Connect for access to services such as Amazon S3?

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

Answer: C

NEW QUESTION 282

Your application is using an ELB in front of an Auto Scaling group of web/application servers deployed across two AZs and a Multi-AZ RDS Instance for data persistence.

The database CPU is often above 80% usage and 90% of I/O operations on the database are reads. To improve performance you recently added a single-node Memcached ElastiCache Cluster to cache frequent DB query results. In the next weeks the overall workload is expected to grow by 30%.

Do you need to change anything in the architecture to maintain the high availability of the application with the anticipated additional load? Why?

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

Answer: A

NEW QUESTION 284

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 286

Auto Scaling requests are signed with a signature calculated from the request and the user's private key.

- A. SSL
- B. AES-256
- C. HMAC-SHA1
- D. X.509

Answer: C

NEW QUESTION 288

What does elasticity mean to AWS?

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

Answer: B

NEW QUESTION 292

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 293

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

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

Answer: B

NEW QUESTION 295

A newspaper organization has a on-premises application which allows the public to search its back catalogue and retrieve individual newspaper pages via a website written in Java They have scanned the old newspapers into JPEGs (approx 17TB) and used Optical Character Recognition (OCR) to populate a commercial search product. The hosting platform and software are now end of life and the organization wants to migrate its archive to AWS and produce a cost efficient architecture and still be designed for availability and durability. Which is the most appropriate?

- A. Use S3 with reduced redundancy to store and serve the scanned files, install the commercial search application on EC2 Instances and configure with auto-scaling and an Elastic Load Balancer.
- B. Model the environment using CloudFormation use an EC2 instance running Apache webserver and an open source search application, stripe multiple standard EBS volumes together to store the JPEGs and search index.
- C. Use S3 with standard redundancy to store and serve the scanned files, use CloudSearch for query processing, and use Elastic Beanstalk to host the website across multiple availability zones.
- D. Use a single-AZ RDS MySQL instance to store the search index and the JPEG images use an EC2 instance to serve the website and translate user queries into SQL.
- E. Use a CloudFront download distribution to serve the JPEGs to the end users and Install the current commercial search product, along with a Java Container on the website on EC2 instances and use Route53 with DNS round-robin.

Answer: C

NEW QUESTION 300

.....

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>