

Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional



NEW QUESTION 1

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 2

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

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

Answer: D

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. AWS provides two features that the user can use to increase security in VPC: security groups and network ACLs. Security group works at the instance level while ACL works at the subnet level. ACL allows both allow and deny rules.

Thus, when the user wants to reject traffic from the selected IPs it is recommended to use ACL with subnets.

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

NEW QUESTION 3

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

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

Answer: D

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. The VPC also work with IAM and the organization can create IAM users who have access to various VPC services. The organization can setup access for the IAM user who can modify the security groups of the VPC. The sample policy is given below:

```
{
  "Version": "2012-10-17",
  "Statement":
  [
    { "Effect": "Allow", "Action": "ec2:RunInstances", "Resource":
    ["arn:aws:ec2:region::image/ami-*", "arn:aws:ec2:region:account:subnet/subnet-1a2b3c4d", "arn:aws:ec2:region:account:network-interface/*",
    "arn:aws:ec2:region:account:volume/*", "arn:aws:ec2:region:account:key-pair/*", "arn:aws:ec2:region:account:security-group/sg-123abc123" ]
    }
  ]
}
```

With this policy the user can create four subnets in separate zones and provide IAM user access to each subnet

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

NEW QUESTION 4

An organization is planning to host an application on the AWS VPC. The organization wants dedicated instances. However, an AWS consultant advised the organization not to use dedicated instances with VPC as the design has a few limitations. Which of the below mentioned statements is not a limitation of dedicated instances with VPC?

- A. All instances launched with this VPC will always be dedicated instances and the user cannot use a default tenancy model for them.
- B. It does not support the AWS RDS with a dedicated tenancy VPC.
- C. The user cannot use Reserved Instances with a dedicated tenancy model.
- D. The EBS volume will not be on the same tenant hardware as the EC2 instance though the user has configured dedicated tenancy.

Answer: C

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. Dedicated instances are Amazon EC2 instances that run in a Virtual Private Cloud (VPC) on hardware that is dedicated to a single customer. The client's dedicated instances are physically isolated at the host hardware level

from instances that are not dedicated instances as well as from instances that belong to other AWS accounts.

All instances launched with the dedicated tenancy model of VPC will always be dedicated instances. Dedicated tenancy has a limitation that it may not support a few services, such as RDS. Even the EBS will not be on dedicated hardware. However the user can save some cost as well as reserve some capacity by using a Reserved Instance model with dedicated tenancy.

Reference: <http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/dedicated-instance.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

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

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

Answer: C

Explanation:

Generally AWS EC2 allows running 20 on-demand instances and 100 spot instances at a time. This limit can be increased by requesting at

<https://aws.amazon.com/contact-us/ec2-request>. Excluding certain types of instances, the limit is lower than mentioned above. For g2.2xlarge, the user can run only 5

on-demand instance at a time.

Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_ec2

NEW QUESTION 7

A user authenticating with Amazon Cognito will go through a multi-step process to bootstrap their credentials. Amazon Cognito has two different flows for authentication with public providers. Which of the following are the two flows?

- A. Authenticated and non-authenticated
- B. Public and private
- C. Enhanced and basic
- D. Single step and multistep

Answer: C

Explanation:

A user authenticating with Amazon Cognito will go through a multi-step process to bootstrap their credentials. Amazon Cognito has two different flows for authentication with public providers: enhanced and basic.

Reference: <http://docs.aws.amazon.com/cognito/devguide/identity/concepts/authentication-flow/>

NEW QUESTION 8

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 9

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 10

An organization (account ID 123412341234) has configured the IAM policy to allow the user to modify his credentials. What will the below mentioned statement allow the user to perform?

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow", "Action": [ "iam:AddUserToGroup",
    "iam:RemoveUserFromGroup", "iam:GetGroup"
    ]
  }],
  "Resource": "arn:aws:iam:: 123412341234:group/TestingGroup"
}
```

- A. Allow the IAM user to update the membership of the group called TestingGroup
- B. The IAM policy will throw an error due to an invalid resource name
- C. The IAM policy will allow the user to subscribe to any IAM group
- D. Allow the IAM user to delete the TestingGroup

Answer: A

Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the organization (account ID 123412341234) wants their users to manage their subscription to the groups, they should create a relevant policy for that. The below mentioned policy allows the respective IAM user to update the membership of the group called MarketingGroup.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow", "Action": [ "iam:AddUserToGroup",
    "iam:RemoveUserFromGroup", "iam:GetGroup"
    ]
  }],
  "Resource": "arn:aws:iam:: 123412341234:group/ TestingGroup "
}
```

Reference:

<http://docs.aws.amazon.com/IAM/latest/UserGuide/Credentials-Permissions-examples.html#creds-policy-credentials>

NEW QUESTION 10

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

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

Answer: B

Explanation:

If the user is not experiencing the expected IOPS or throughput that is provisioned, ensure that the EC2 bandwidth is not the limiting factor, the instance is EBS-optimized (or include 10 Gigabit network connectMty) and the instance type EBS dedicated bandwidth exceeds the IOPS more than he has provisioned.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-io-characteristics.html>

NEW QUESTION 13

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 actMties.

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

NEW QUESTION 17

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 22

You want to define permissions for a role in an IAM policy. Which of the following configuration formats should you use?

- A. An XML document written in the IAM Policy Language
- B. An XML document written in a language of your choice
- C. A JSON document written in the IAM Policy Language
- D. A JSON document written in a language of your choice

Answer: C

Explanation:

You define the permissions for a role in an IAM policy. An IAM policy is a JSON document written in the IAM Policy Language.

Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_terms-and-concepts.html

NEW QUESTION 24

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

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

Answer: D

Explanation:

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

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

Reference:

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

NEW QUESTION 29

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

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

Answer: C

Explanation:

AWS provides an on demand, scalable infrastructure. AWS EC2 allows the user to launch On-Demand instances and the organization should create an AMI of the running instance. Copy the AMI to another region to enable Disaster Recovery (DR) in case of region failure. The organization should also use EBS for persistent storage and take a snapshot every 10 minutes to meet Recovery time objective (RTO). They should also setup an elastic IP and use it with Route 53 to route requests to the same IP.

When one of the instances fails the organization can launch new instances and assign the same EIP to a new instance to achieve High Availability (HA). The ELB works only for a particular region and does not route requests across regions.

Reference: http://d36cz9buwru1tt.c|oudfront.net/AWS_Disaster_Recovery.pdf

NEW QUESTION 33

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

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

Answer: A

Explanation:

In Amazon ElastiCache, you can run a maximum of 20 Cache Nodes. Reference: <http://aws.amazon.com/elasticache/faqs/>

NEW QUESTION 34

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

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

Answer: D

Explanation:

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

NEW QUESTION 38

ExamKiller has three separate departments and each department has their own AWS accounts. The HR department has created a file sharing site where all the on roll employees' data is uploaded. The Admin department uploads data about the employee presence in the office to their DB hosted in the VPC. The Finance department needs to access data from the HR department to know the on roll employees to calculate the salary based on the number of days that an employee is present in the office.

How can ExamKiller setup this scenario?

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

Answer: C

Explanation:

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

Reference:

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

NEW QUESTION 41

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 44

In Amazon Cognito, your mobile app authenticates with the Identity Provider (IdP) using the provider's SDK. Once the end user is authenticated with the IdP, the OAuth or OpenID Connect token returned from the IdP is passed by your app to Amazon Cognito, which returns a new for the user and a set of temporary, limited-prMlege AWS credentials.

- A. Cognito Key Pair
- B. Cognito API
- C. Cognito ID
- D. Cognito SDK

Answer: C

Explanation:

Your mobile app authenticates with the identity provider (IdP) using the provider's SDK. Once the end user is authenticated with the IdP, the OAuth or OpenID

Connect token returned from the IdP is passed by your app to Amazon Cognito, which returns a new Cognito ID for the user and a set of temporary, limited-prMlege AWS credentials.
Reference: <http://aws.amazon.com/cognito/faqs/>

NEW QUESTION 48

What is the maximum length for a certificate ID in AWS IAM?

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

Answer: D

Explanation:

The maximum length for a certificate ID is 128 characters.
Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html>

NEW QUESTION 50

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_pio ps

NEW QUESTION 53

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

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

Answer: C

Explanation:

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

NEW QUESTION 55

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 59

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 63

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 68

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

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

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

Answer: A

Explanation:

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

Reference:

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

NEW QUESTION 69

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 74

An organization is making software for the CIA in US

- A. CIA agreed to host the application on AWS but in a secure environment
- B. The organization is thinking of hosting the application on the AWS GovCloud region
- C. Which of the below mentioned difference is not correct when the organization is hosting on the AWS GovCloud in comparison with the AWS standard region?
- D. The billing for the AWS GovCloud will be in a different account than the Standard AWS account.
- E. GovCloud region authentication is isolated from Amazon.com.
- F. Physical and logical administrative access only to U.S. persons.
- G. persons.
- H. It is physically isolated and has logical network isolation from all the other region

Answer: A

Explanation:

AWS GovCloud (US) is an isolated AWS region designed to allow U.S. government agencies and customers to move sensitive workloads into the cloud by addressing their specific regulatory and compliance requirements. The AWS GovCloud (US) Region adheres to the U.S. International Traffic in Arms Regulations (ITAR) requirements. It has added advantages, such as: Restricting physical and logical administrative access to U.S. persons only. There will be a separate AWS GovCloud (US) credentials, such as access key and secret access key than the standard AWS account. The user signs in with the IAM user name and password. The AWS GovCloud (US) Region authentication is completely isolated from Amazon.com. If the organization is planning to host on EC2 in AWS GovCloud then it will be billed to standard AWS account of organization since AWS GovCloud billing is linked with the standard AWS account and is not be billed separately. Reference: <http://docs.aws.amazon.com/govcloud-us/latest/UserGuide/whatis.html>

NEW QUESTION 75

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

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

Answer: C

Explanation:

Provisioned IOPS volumes are designed to meet the needs of I/O-intensive workloads, particularly database workloads that are sensitive to storage performance and consistency in random access I/O throughput. Provisioned IOPS volumes are designed to meet the needs of I/O-intensive workloads, particularly database workloads, that are sensitive to storage performance and consistency in random access I/O throughput business applications, database workloads, such as NoSQL DB, RDBMS, etc. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html>

NEW QUESTION 79

In the context of policies and permissions in AWS IAM, the Condition element is .

- A. crucial while writing the IAM policies
- B. an optional element
- C. always set to null
- D. a mandatory element

Answer: B

Explanation:

The Condition element (or Condition block) lets you specify conditions for when a policy is in effect. The Condition element is optional. Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html

NEW QUESTION 84

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

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

Answer: A

Explanation:

When using string conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions. For instance, streq is the short version of StringEqualsIgnoreCase that checks for the exact match between two strings ignoring their case. Reference: <http://awsdocs.s3.amazonaws.com/SNS/20100331/sns-gsg-2010-03-31.pdf>

NEW QUESTION 88

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) environment.

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

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

- A. higher latency and lower throughput.
- B. lower latency and higher throughput.

- C. higher throughput only.
- D. higher latency onl

Answer: B

Explanation:

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

NEW QUESTION 95

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 96

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

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

Answer: C

Explanation:

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

NEW QUESTION 98

Within an IAM policy, can you add an IfExists condition at the end of a Null condition?

- A. Yes, you can add an IfExists condition at the end of a Null condition but not in all Regions.
- B. Yes, you can add an IfExists condition at the end of a Null condition depending on the condition.
- C. No, you cannot add an IfExists condition at the end of a Null condition.
- D. Yes, you can add an IfExists condition at the end of a Null conditio

Answer: C

Explanation:

Within an IAM policy, IfExists can be added to the end of any condition operator except the Null condition. It can be used to indicate that conditional comparison needs to happen if the policy key is present in the context of a request; otherwise, it can be ignored.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html

NEW QUESTION 101

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

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

Answer: A

Explanation:

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

key, while the user is given the private key.

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

NEW QUESTION 105

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 110

What happens when Dedicated instances are launched into a VPC?

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

Answer: C

Explanation:

If you launch an instance into a VPC that has an instance tenancy of dedicated, your instance is automatically a Dedicated instance, regardless of the tenancy of the instance.

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

NEW QUESTION 112

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

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

Answer: A

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources, such as RDS into a virtual network that the user has defined. Subnets are segments of a VPC's IP address range that the user can designate to a group of VPC resources based on security and operational needs. A DB subnet group is a collection of subnets (generally private) that the user can create in a VPC and assign to the RDS DB instances. A DB subnet group allows the user to specify a particular VPC when creating the DB instances.

Each DB subnet group should have subnets in at least two Availability Zones in a given region. If the RDS instance is required to be accessible from the internet the organization must enable the VPC attributes, DNS hostnames and DNS resolution. For each RDS DB instance that the user runs in a VPC, he should reserve at least one address in each subnet in the DB subnet group for use by Amazon RDS for recovery actions.

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

NEW QUESTION 116

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 119

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

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

Answer: D

Explanation:

AWS Direct Connect is a network service that provides an alternative to using the Internet to utilize AWS cloud service. Hence, a Direct Connect link cannot be connected to the Internet directly.

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

NEW QUESTION 122

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 VP
- 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-acces>

NEW QUESTION 125

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 128

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 129

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

- A. Raid 1
- B. Raid 0

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

Answer: C

Explanation:

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

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

NEW QUESTION 130

In the context of IAM roles for Amazon EC2, which of the following NOT true about delegating permission to make API requests?

- A. You cannot create an IAM role.
- B. You can have the application retrieve a set of temporary credentials and use them.
- C. You can specify the role when you launch your instances.
- D. You can define which accounts or AWS services can assume the rol

Answer: A

Explanation:

Amazon designed IAM roles so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that the applications use. Instead of creating and distributing your AWS credentials, you can delegate permission to make API requests using IAM roles as follows: Create an IAM role. Define which accounts or AWS services can assume the role. Define which API actions and resources the application can use after assuming the role. Specify the role when you launch your instances. Have the application retrieve a set of temporary credentials and use them.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

NEW QUESTION 135

In Amazon Cognito what is a silent push notification?

- A. It is a push message that is received by your application on a user's device that will not be seen by the user
- B. It is a push message that is received by your application on a user's device that will return the user's geolocation.
- C. It is a push message that is received by your application on a user's device that will not be heard by the user
- D. It is a push message that is received by your application on a user's device that will return the user's authentication credentials.

Answer: A

Explanation:

Amazon Cognito uses the Amazon Simple Notification Service (SNS) to send silent push notifications to devices. A silent push notification is a push message that is received by your application on a user's device that will not be seen by the user.

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

NEW QUESTION 137

When using Numeric Conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions. Which of the following is the short version of the Numeric Condition "NumericLessThanEquals"?

- A. numlteq
- B. numlteql
- C. numltequals
- D. numeqql

Answer: A

Explanation:

When using Numeric Conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions. For instance, numlteq is the short version of NumericLessThanEquals.

Reference: <http://awsdocs.s3.amazonaws.com/SQS/2011-10-01/sqs-dg-2011-10-01.pdf>

NEW QUESTION 141

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

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

Answer: A

Explanation:

To enable running actMties using on-premise resources, AWS Data Pipeline does the following: It supply a Task Runner package that can be installed on your on-premise hosts.

This package continuously polls the AWS Data Pipeline service for work to perform.

When it's time to run a particular actMty on your on-premise resources, it will issue the appropriate command to the Task Runner.

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

NEW QUESTION 146

Once the user has set ElastiCache for an application and it is up and running, which services, does Amazon not provide for the user:

- A. The ability for client programs to automatically identify all of the nodes in a cache cluster, and to initiate and maintain connections to all of these nodes
- B. Automating common administrative tasks such as failure detection and recovery, and software patching
- C. Providing default Time To Live (TTL) in the AWS ElastiCache Redis Implementation for different type of data.
- D. Providing detailed monitoring metrics associated with your Cache Nodes, enabling you to diagnose and react to issues very quickly

Answer: C

Explanation:

Amazon provides failure detection and recovery, and software patching and monitoring tools which is called CloudWatch. In addition it provides also Auto Discovery to automatically identify and initialize all nodes of cache cluster for Amazon ElastiCache.
Reference: <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/WhatIs.html>

NEW QUESTION 149

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

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

Answer: A

Explanation:

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

NEW QUESTION 154

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 159

A user is planning to host a Highly Available system on the AWS VPC. Which of the below mentioned statements is helpful in this scenario?

- A. Create VPC subnets in two separate availability zones and launch instances in different subnets.
- B. Create VPC with only one public subnet and launch instances in different AZs using that subnet.
- C. Create two VPCs in two separate zones and setup failover with ELB such that if one VPC fails it will divert traffic to another VPC.
- D. Create VPC with only one private subnet and launch instances in different AZs using that subne

Answer: A

Explanation:

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

NEW QUESTION 163

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 164

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 168

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 173

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

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

Answer: C

Explanation:

In AWS Data Pipeline, you can define Amazon SNS alarms to trigger on actMties such as success, failure, or delay by creating an alarm object and referencing it in the onFail, onSuccess, or onLate slots of the actMty object.

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

NEW QUESTION 178

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

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

Answer: D

Explanation:

For a single-node configuration in Amazon Redshift, the only option available is the smallest of the two options. The 8XL extra-large nodes are only available in a multi-node configuration

Reference: <http://docs.aws.amazon.com/redshift/latest/mgmt/working-with-clusters.html>

NEW QUESTION 180

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

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

Answer: B

Explanation:

The basic building block of Amazon RDS is the DB instance. DB instance storage comes in three types: Magnetic, General Purpose (SSD), and Provisioned IOPS

(SSD). When you buy a server, you get CPU, memory, storage, and IOPS, all bundled together. With Amazon RDS, these are split apart so that you can scale them independently. So, for example, if you need more CPU, less IOPS, or more storage, you can easily allocate them.

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

NEW QUESTION 181

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 184

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 188

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 192

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

- A. <https://Q99988887777.aws.amazon.com/examkiller/>
- B. <https://signin.aws.amazon.com/examkiller/>
- C. <https://examkiller.signin.aws.amazon.com/Q99988887777/console/>
- D. <https://Q99988887777.signin.aws.amazon.com/console/>

Answer: D

Explanation:

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

Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/AccountAlias.html>

NEW QUESTION 195

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 IAM 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 199

To serve Web traffic for a popular product your chief financial officer and IT director have purchased 10 m1 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 m1 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 instances.
- 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 200

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 them.
- E. Use SES to send emails to customers.

Answer: C

NEW QUESTION 203

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 to 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 CloudWatch.
- 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 206

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

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

N. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.

Answer: D

NEW QUESTION 208

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

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

Answer: ADF

NEW QUESTION 211

You have been asked to design the storage layer for an application. The application requires disk performance of at least 100,000 IOPS. In addition, the storage layer must be able to survive the loss of an individual disk, EC2 instance, or Availability Zone without any data loss. The volume you provide must have a capacity of at least 3 TB. Which of the following designs will meet these objectives?

- A. Instantiate a c3.8xlarge instance in us-east-1. Provision 4x1TB EBS volumes, attach them to the instance, and configure them as a single RAID 5 volume.
- B. Ensure that EBS snapshots are performed every 15 minutes.
- C. Instantiate a c3.8xlarge instance in us-east-1. Provision 3x1TB EBS volumes, attach them to the instance, and configure them as a single RAID 0 volume.
- D. Ensure that EBS snapshots are performed every 15 minutes.
- E. Instantiate an i2.8xlarge instance in us-east-1.
- F. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instance.
- G. Provision 3x1TB EBS volumes, attach them to the instance, and configure them as a second RAID 0 volume.
- H. Configure synchronous, block-level replication from the ephemeral-backed volume to the EBS-backed volume.
- I. Instantiate a c3.8xlarge instance in us-east-1. Provision an AWS Storage Gateway and configure it for 3 TB of storage and 100,000 IOP.
- J. Attach the volume to the instance.
- K. Instantiate an i2.8xlarge instance in us-east-1.
- L. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instance.
- M. Configure synchronous, block-level replication to an identically configured instance in us-east-1b.

Answer: C

NEW QUESTION 212

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 215

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

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

Answer: A

NEW QUESTION 219

An AWS customer is deploying an application that is composed of an AutoScaling group of EC2 instances. The customer's security policy requires that every outbound connection from these instances to any other service within the customer's Virtual Private Cloud must be authenticated using a unique x.509 certificate that contains the specific instance-id.

In addition an x 509 certificates must Designed by the customer's Key management service in order to be trusted for authentication. Which of the following configurations will support these requirements?

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

Answer: A

NEW QUESTION 222

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

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

Answer: A

NEW QUESTION 225

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

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

Answer: A

NEW QUESTION 230

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

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

Answer: A

NEW QUESTION 233

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

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

Answer: C

NEW QUESTION 237

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 240

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 243

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 246

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 247

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 day
- C. CloudFront to serve HLS transcoded videos from EC2.
- D. Elastic Transcoder to transcode original high-resolution MP4 videos to HL
- E. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few day
- F. CloudFront to serve HLS transcoded videos from EC2.
- G. Elastic Transcoder to transcode original high-resolution MP4 videos to HL
- H. S3 to host videos with Lifecycle Management to archive original files to Glacier after a few day
- I. C|oudFront 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 251

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

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

Answer: AC

NEW QUESTION 252

You are running a news website in the eu-west-1 region that updates every 15 minutes. The website has a world-wide audience it uses an Auto Scaling group behind an Elastic Load Balancer and an Amazon RDS database. Static content resides on Amazon S3, and is distributed through Amazon CloudFront. Your Auto Scaling group is set to trigger a scale up event at 60% CPU utilization, you use an Amazon RDS extra large DB instance with 10,000 Provisioned IOPS its CPU utilization is around 80%. While freeable memory is in the 2 GB range.

Web analytics reports show that the average load time of your web pages is around 1.5 to 2 seconds, but your SEO consultant wants to bring down the average load time to under 0.5 seconds.

How would you improve page load times for your users? (Choose 3 answers)

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

Answer: ABD

NEW QUESTION 253

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

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

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

Answer: C

NEW QUESTION 256

You are designing an S3UTLS solution that requires HTTPS clients to be authenticated by the Webserver using client certificate authentication. The solution must be resilient.

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

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

Answer: AB

NEW QUESTION 260

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 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 users 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 users 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 users vote and store the result into a DynamoDB table using IAM Roles for EC2 instances to gain permissions to the DynamoDB table.
- D. Use 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 users 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 264

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

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

Answer: B

NEW QUESTION 265

Your fortune 500 company has undertaken 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 267

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

How would you optimize this scenario to solve performance issues and automate the process as much as possible?

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

Answer: A

NEW QUESTION 269

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

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

Answer: C

NEW QUESTION 270

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

Which AWS storage and database architecture meets the requirements of the application?

- A. Web servers: store read-only data in S3, and copy from S3 to root volume at boot time
- B. App servers: share state using a combination of DynamoDB and IP unicast
- C. Database: use RDS with multi-AZ deployment and one or more read replicas

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

Answer: C

NEW QUESTION 275

Your company plans to host a large donation website on Amazon Web Services (AWS). You anticipate a large and undetermined amount of traffic that will create many database writes. To be certain that you do not drop any writes to a database hosted on AWS. Which service should you use?

- A. Amazon RDS with provisioned IOPS up to the anticipated peak write throughput.
- B. Amazon Simple Queue Service (SQS) for capturing the writes and draining the queue to write to the database.
- C. Amazon ElastiCache to store the writes until the writes are committed to the database.
- D. Amazon DynamoDB with provisioned write throughput up to the anticipated peak write throughput

Answer: B

NEW QUESTION 279

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

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

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

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

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

rtb-218bc449 rtb-238bc44b Associations:

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

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

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

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

Answer: A

NEW QUESTION 281

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 CloudFormation 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 283

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 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 288

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 290

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 291

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

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

Answer: C

NEW QUESTION 294

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

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

Answer: D

NEW QUESTION 299

You've been hired to enhance the overall security posture for a very large e-commerce site. They have a well architected multi-tier application running in a VPC that uses ELBs in front of both the web and the app tier with static assets served directly from S3. They are using a combination of RDS and DynamoDB for their dynamic data and then archive nightly into S3 for further processing with EMR. They are concerned because they found questionable log entries and suspect someone is attempting to gain unauthorized access.

Which approach provides a cost effective scalable mitigation to this kind of attack?

- A. Recommend that they lease space at a DirectConnect partner location and establish a 1G DirectConnect connection to their VPC. They would then establish Internet connectivity into their space, filter the traffic in hardware Web Application Firewall (WAF). And then pass the traffic through the DirectConnect connection into their application running in their VPC.
- B. Add previously identified hostile source IPs as an explicit INBOUND DENY NACL to the web tier subnet
- C. Add a WAF tier by creating a new ELB and an AutoScaling group of EC2 instances running a host-based WAF. They would redirect Route 53 to resolve to the new WAF tier ELB. The WAF tier would then pass the traffic to the current web tier. The web tier Security Groups would be updated to only allow traffic from the WAF tier Security Group
- D. Remove all but TLS 1.2 from the web tier ELB and enable Advanced Protocol Filtering. This will enable the ELB itself to perform WAF functionality.

Answer: C

NEW QUESTION 301

Your team has a tomcat-based Java application you need to deploy into development, test and production environments. After some research, you opt to use Elastic Beanstalk due to its tight integration with your developer tools and RDS due to its ease of management. Your QA team lead points out that you need to roll a sanitized set of production data into your environment on a nightly basis.

Similarly, other software teams in your org want access to that same restored data via their EC2 instances in your VPC. The optimal setup for persistence and security that meets the above requirements would be the following.

- A. Create your RDS instance as part of your Elastic Beanstalk definition and alter its security group to allow access to it from hosts in your application subnets.

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

Answer: A

NEW QUESTION 304

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 308

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

What DR strategy could be used to achieve this RTO and RPO in the event of this kind of failure?

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

Answer: A

NEW QUESTION 311

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

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

Answer: ABC

NEW QUESTION 314

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 319

The following are AWS Storage services? Choose 2 answers.

- A. AWS Relational Database Service (AWS RDS)
- B. AWS ElastiCache
- C. AWS Glacier
- D. AWS Import/Export

Answer: BD

NEW QUESTION 321

How is AWS readily distinguished from other vendors in the traditional IT computing landscape?

- A. Experience

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

Answer: C

NEW QUESTION 324

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 326

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

AWS-Certified-Solutions-Architect-Professional Practice Exam Features:

- * AWS-Certified-Solutions-Architect-Professional Questions and Answers Updated Frequently
- * AWS-Certified-Solutions-Architect-Professional Practice Questions Verified by Expert Senior Certified Staff
- * AWS-Certified-Solutions-Architect-Professional Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * AWS-Certified-Solutions-Architect-Professional Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-Certified-Solutions-Architect-Professional Practice Test Here](#)