

AWS-Solution-Architect-Associate Dumps

Amazon AWS Certified Solutions Architect - Associate

<https://www.certleader.com/AWS-Solution-Architect-Associate-dumps.html>



NEW QUESTION 1

You are trying to launch an EC2 instance, however the instance seems to go into a terminated status immediately. What would probably not be a reason that this is happening?

- A. The AMI is missing a required part.
- B. The snapshot is corrupt.
- C. You need to create storage in EBS first.
- D. You've reached your volume limit

Answer: C

Explanation:

Amazon EC2 provides a virtual computing environments, known as an instance.

After you launch an instance, AWS recommends that you check its status to confirm that it goes from the pending status to the running status, the not terminated status.

The following are a few reasons why an Amazon EBS-backed instance might immediately terminate: You've reached your volume limit.

The AMI is missing a required part. The snapshot is corrupt. Reference:

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_InstanceStraightToTerminated.html

NEW QUESTION 2

Can a user get a notification of each instance start / terminate configured with Auto Scaling?

- A. Yes, if configured with the Launch Config
- B. Yes, always
- C. Yes, if configured with the Auto Scaling group
- D. No

Answer: C

Explanation:

The user can get notifications using SNS if he has configured the notifications while creating the Auto Scaling group.

Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/GettingStartedTutorial.html>

NEW QUESTION 3

You are migrating an internal server on your DC to an EC2 instance with EBS volume. Your server disk usage is around 500GB so you just copied all your data to a 2TB disk to be used with AWS Import/Export. Where will the data be imported once it arrives at Amazon?

- A. to a 2TB EBS volume
- B. to an S3 bucket with 2 objects of 1TB
- C. to an 500GB EBS volume
- D. to an S3 bucket as a 2TB snapshot

Answer: B

Explanation:

An import to Amazon EBS will have different results depending on whether the capacity of your storage device is less than or equal to 1 TB or greater than 1 TB.

The maximum size of an Amazon EBS snapshot is 1 TB, so if the device image is larger than 1 TB, the image is chunked and stored on Amazon S3. The target location is determined based on the total capacity of the device, not the amount of data on the device.

Reference: <http://docs.aws.amazon.com/AWSImportExport/latest/DG/Concepts.html>

NEW QUESTION 4

True or false? A VPC contains multiple subnets, where each subnet can span multiple Availability Zones.

- A. This is true only if requested during the set-up of VPC.
- B. This is true.
- C. This is false.
- D. This is true only for US region

Answer: C

Explanation:

A VPC can span several Availability Zones. In contrast, a subnet must reside within a single Availability Zone.

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

NEW QUESTION 5

An edge location refers to which Amazon Web Service?

- A. An edge location is referred to the network configured within a Zone or Region
- B. An edge location is an AWS Region
- C. An edge location is the location of the data center used for Amazon CloudFront.
- D. An edge location is a Zone within an AWS Region

Answer: C

Explanation:

Amazon CloudFront is a content distribution network. A content delivery network or content distribution network (CDN) is a large distributed system of servers

deployed in multiple data centers across the world. The location of the data center used for CDN is called edge location. Amazon CloudFront can cache static content at each edge location. This means that your popular static content (e.g., your site's logo, navigational images, cascading style sheets, JavaScript code, etc.) will be available at a nearby edge location for the browsers to download with low latency and improved performance for viewers. Caching popular static content with Amazon CloudFront also helps you offload requests for such files from your origin server — CloudFront serves the cached copy when available and only makes a request to your origin server if the edge location receiving the browser's request does not have a copy of the file. Reference: <http://aws.amazon.com/cloudfront/>

NEW QUESTION 6

You are looking at ways to improve some existing infrastructure as it seems a lot of engineering resources are being taken up with basic management and monitoring tasks and the costs seem to be excessive.

You are thinking of deploying Amazon ElastiCache to help. Which of the following statements is true in regards to ElastiCache?

- A. You can improve load and response times to user actions and queries however the cost associated with scaling web applications will be more.
- B. You can't improve load and response times to user actions and queries but you can reduce the cost associated with scaling web applications.
- C. You can improve load and response times to user actions and queries however the cost associated with scaling web applications will remain the same.
- D. You can improve load and response times to user actions and queries and also reduce the cost associated with scaling web applications.

Answer: D

Explanation:

Amazon ElastiCache is a web service that makes it easy to deploy and run Memcached or Redis protocol-compliant server nodes in the cloud. Amazon ElastiCache improves the performance of web applications by allowing you to retrieve information from a fast, managed, in-memory caching system, instead of relying entirely on slower disk-based databases. The service simplifies and offloads the management, monitoring and operation of in-memory cache environments, enabling your engineering resources to focus on developing applications.

Using Amazon ElastiCache, you can not only improve load and response times to user actions and queries, but also reduce the cost associated with scaling web applications.

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

NEW QUESTION 7

Do Amazon EBS volumes persist independently from the running life of an Amazon EC2 instance?

- A. Yes, they do but only if they are detached from the instance.
- B. No, you cannot attach EBS volumes to an instance.
- C. No, they are dependent.
- D. Yes, they do

Answer: D

Explanation:

An Amazon EBS volume behaves like a raw, unformatted, external block device that you can attach to a single instance. The volume persists independently from the running life of an Amazon EC2 instance. Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Storage.html>

NEW QUESTION 8

You need to migrate a large amount of data into the cloud that you have stored on a hard disk and you decide that the best way to accomplish this is with AWS Import/Export and you mail the hard disk to AWS. Which of the following statements is incorrect in regards to AWS Import/Export?

- A. It can export from Amazon S3
- B. It can Import to Amazon Glacier
- C. It can export from Amazon Glacier.
- D. It can Import to Amazon EBS

Answer: C

Explanation:

AWS Import/Export supports: Import to Amazon S3

Export from Amazon S3 Import to Amazon EBS Import to Amazon Glacier

AWS Import/Export does not currently support export from Amazon EBS or Amazon Glacier. Reference:

<https://docs.aws.amazon.com/AWSImportExport/latest/DG/whatdisk.html>

NEW QUESTION 9

In Amazon EC2, partial instance-hours are billed .

- A. per second used in the hour
- B. per minute used
- C. by combining partial segments into full hours
- D. as full hours

Answer: D

Explanation:

Partial instance-hours are billed to the next hour. Reference: <http://aws.amazon.com/ec2/faqs/>

NEW QUESTION 10

Can you specify the security group that you created for a VPC when you launch an instance in EC2-Classic?

- A. No, you can specify the security group created for EC2-Classic when you launch a VPC instance.
- B. No

- C. Yes
- D. No, you can specify the security group created for EC2-Classic to a non-VPC based instance onl

Answer: B

Explanation:

If you're using EC2-Classic, you must use security groups created specifically for EC2-Classic. When you launch an instance in EC2-Classic, you must specify a security group in the same region as the instance. You can't specify a security group that you created for a VPC when you launch an instance in EC2-Classic.

Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html#ec2-classic-security-groups>

NEW QUESTION 10

You have been asked to build a database warehouse using Amazon Redshift. You know a little about it, including that it is a SQL data warehouse solution, and uses industry standard ODBC and JDBC connections and PostgreSQL drivers. However you are not sure about what sort of storage it uses for database tables. What sort of storage does Amazon Redshift use for database tables?

- A. InnoDB Tables
- B. NDB data storage
- C. Columnar data storage
- D. NDB CLUSTER Storage

Answer: C

Explanation:

Amazon Redshift achieves efficient storage and optimum query performance through a combination of massively parallel processing, columnar data storage, and very efficient, targeted data compression encoding schemes.

Columnar storage for database tables is an important factor in optimizing analytic query performance because it drastically reduces the overall disk I/O requirements and reduces the amount of data you need to load from disk.

Reference: http://docs.aws.amazon.com/redshift/latest/dg/c_columnar_storage_disk_mem_mgmt.html

NEW QUESTION 15

You are checking the workload on some of your General Purpose (SSD) and Provisioned IOPS (SSD) volumes and it seems that the I/O latency is higher than you require. You should probably check the to make sure that your application is not trying to drive more IOPS than you have provisioned.

- A. Amount of IOPS that are available
- B. Acknowledgement from the storage subsystem
- C. Average queue length
- D. Time it takes for the I/O operation to complete

Answer: C

Explanation:

In EBS workload demand plays an important role in getting the most out of your General Purpose (SSD) and Provisioned IOPS (SSD) volumes. In order for your volumes to deliver the amount of IOPS that are available, they need to have enough I/O requests sent to them. There is a relationship between the demand on the volumes, the amount of IOPS that are available to them, and the latency of the request (the amount of time it takes for the I/O operation to complete).

Latency is the true end-to-end client time of an I/O operation; in other words, when the client sends a IO, how long does it take to get an acknowledgement from the storage subsystem that the IO read or write is complete.

If your I/O latency is higher than you require, check your average queue length to make sure that your application is not trying to drive more IOPS than you have provisioned. You can maintain high IOPS while keeping latency down by maintaining a low average queue length (which is achieved by provisioning more IOPS for your volume).

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

NEW QUESTION 19

You have been given a scope to deploy some AWS infrastructure for a large organisation. The requirements are that you will have a lot of EC2 instances but may need to add more when the average utilization of your Amazon EC2 fleet is high and conversely remove them when CPU utilization is low. Which AWS services would be best to use to accomplish this?

- A. Auto Scaling, Amazon CloudWatch and AWS Elastic Beanstalk
- B. Auto Scaling, Amazon CloudWatch and Elastic Load Balancing.
- C. Amazon CloudFront, Amazon CloudWatch and Elastic Load Balancing.
- D. AWS Elastic Beanstalk , Amazon CloudWatch and Elastic Load Balancing

Answer: B

Explanation:

Auto Scaling enables you to follow the demand curve for your applications closely, reducing the need to manually provision Amazon EC2 capacity in advance. For example, you can set a condition to add new

Amazon EC2 instances in increments to the Auto Scaling group when the average utilization of your Amazon EC2 fleet is high; and similarly, you can set a condition to remove instances in the same increments when CPU utilization is low. If you have predictable load changes, you can set a schedule through Auto Scaling to plan your scaling activities. You can use Amazon CloudWatch to send alarms to trigger scaling activities and Elastic Load Balancing to help distribute traffic to your instances within Auto Scaling groups. Auto Scaling enables you to run your Amazon EC2 fleet at optimal utilization. Reference:

<http://aws.amazon.com/autoscaling/>

NEW QUESTION 22

Much of your company's data does not need to be accessed often, and can take several hours for retrieval time, so it's stored on Amazon Glacier. However someone within your organization has expressed concerns that his data is more sensitive than the other data, and is wondering whether the high level of encryption that he knows is on S3 is also used on the much cheaper Glacier service. Which of the following statements would be most applicable in

regards to this concern?

- A. There is no encryption on Amazon Glacier, that's why it is cheaper.
- B. Amazon Glacier automatically encrypts the data using AES-128 a lesser encryption method than Amazon S3 but you can change it to AES-256 if you are willing to pay more.
- C. Amazon Glacier automatically encrypts the data using AES-256, the same as Amazon S3.
- D. Amazon Glacier automatically encrypts the data using AES-128 a lesser encryption method than Amazon S3.

Answer: C

Explanation:

Like Amazon S3, the Amazon Glacier service provides low-cost, secure, and durable storage. But where S3 is designed for rapid retrieval, Glacier is meant to be used as an archival service for data that is not accessed often, and for which retrieval times of several hours are suitable. Amazon Glacier automatically encrypts the data using AES-256 and stores it durably in an immutable form. Amazon Glacier is designed to provide average annual durability of 99.999999999% for an archive. It stores each archive in multiple facilities and multiple devices. Unlike traditional systems which can require laborious data verification and manual repair, Glacier performs regular, systematic data integrity checks, and is built to be automatically self-healing. Reference: <http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf>

NEW QUESTION 27

You need to set up a complex network infrastructure for your organization that will be reasonably easy to deploy, replicate, control, and track changes on. Which AWS service would be best to use to help you accomplish this?

- A. AWS Import/Export
- B. AWS CloudFormation
- C. Amazon Route 53
- D. Amazon CloudWatch

Answer: B

Explanation:

AWS CloudFormation is a service that helps you model and set up your Amazon Web Services resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS. You create a template that describes all the AWS resources that you want (like Amazon EC2 instances or Amazon RDS DB instances), and AWS CloudFormation takes care of provisioning and configuring those resources for you. You don't need to manually create and configure AWS resources and figure out what's dependent on what. AWS CloudFormation handles all of that. Reference: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/Welcome.html>

NEW QUESTION 28

In Amazon EC2, while sharing an Amazon EBS snapshot, can the snapshots with AWS Marketplace product codes be public?

- A. Yes, but only for US-based providers.
- B. Yes, they can be public.
- C. No, they cannot be made public.
- D. Yes, they are automatically made public by the system.

Answer: C

Explanation:

Snapshots with AWS Marketplace product codes can't be made public. Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/ebs-modifying-snapshot-permissions.html>

NEW QUESTION 32

Can resource record sets in a hosted zone have a different domain suffix (for example, www.blog.acme.com and www.acme.ca)?

- A. Yes, it can have for a maximum of three different TLDs.
- B. Yes
- C. Yes, it can have depending on the TLD.
- D. No

Answer: D

Explanation:

The resource record sets contained in a hosted zone must share the same suffix. For example, the example.com hosted zone can contain resource record sets for www.example.com and www.aws.example.com subdomains, but it cannot contain resource record sets for a www.example.ca subdomain. Reference: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/AboutHostedZones.html>

NEW QUESTION 33

A user has launched 10 EC2 instances inside a placement group. Which of the below mentioned statements is true with respect to the placement group?

- A. All instances must be in the same AZ
- B. All instances can be across multiple regions
- C. The placement group cannot have more than 5 instances
- D. All instances must be in the same region

Answer: A

Explanation:

A placement group is a logical grouping of EC2 instances within a single Availability Zone. Using placement groups enables applications to participate in a low-

latency, 10 Gbps network. Placement groups are recommended for applications that benefit from low network latency, high network throughput or both.
Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

NEW QUESTION 34

You have been storing massive amounts of data on Amazon Glacier for the past 2 years and now start to wonder if there are any limitations on this. What is the correct answer to your QUESTION ?

- A. The total volume of data is limited but the number of archives you can store are unlimited.
- B. The total volume of data is unlimited but the number of archives you can store are limited.
- C. The total volume of data and number of archives you can store are unlimited.
- D. The total volume of data is limited and the number of archives you can store are limited.

Answer: C

Explanation:

An archive is a durably stored block of information. You store your data in Amazon Glacier as archives. You may upload a single file as an archive, but your costs will be lower if you aggregate your data. TAR and ZIP are common formats that customers use to aggregate multiple files into a single file before uploading to Amazon Glacier.

The total volume of data and number of archives you can store are unlimited. Individual Amazon Glacier archives can range in size from 1 byte to 40 terabytes. The largest archive that can be uploaded in a single upload request is 4 gigabytes.

For items larger than 100 megabytes, customers should consider using the Multipart upload capability. Archives stored in Amazon Glacier are immutable, i.e. archives can be uploaded and deleted but cannot be edited or overwritten.

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

NEW QUESTION 35

Which of the following statements is true of creating a launch configuration using an EC2 instance?

- A. The launch configuration can be created only using the Query APIs.
- B. Auto Scaling automatically creates a launch configuration directly from an EC2 instance.
- C. A user should manually create a launch configuration before creating an Auto Scaling group.
- D. The launch configuration should be created manually from the AWS CL

Answer: B

Explanation:

You can create an Auto Scaling group directly from an EC2 instance. When you use this feature, Auto Scaling automatically creates a launch configuration for you as well.

Reference:

<http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/create-lc-with-instanceID.html>

NEW QUESTION 36

You need to set up a high level of security for an Amazon Relational Database Service (RDS) you have just built in order to protect the confidential information stored in it. What are all the possible security groups that RDS uses?

- A. DB security groups, VPC security groups, and EC2 security groups.
- B. DB security groups only.
- C. EC2 security groups only.
- D. VPC security groups, and EC2 security group

Answer: A

Explanation:

A security group controls the access to a DB instance. It does so by allowing access to IP address ranges or Amazon EC2 instances that you specify.

Amazon RDS uses DB security groups, VPC security groups, and EC2 security groups. In simple terms, a DB security group controls access to a DB instance that is not in a VPC, a VPC security group controls access to a DB instance inside a VPC, and an Amazon EC2 security group controls access to an EC2 instance and can be used with a DB instance.

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

NEW QUESTION 37

After setting up several database instances in Amazon Relational Database Service (Amazon RDS) you decide that you need to track the performance and health of your databases. How can you do this?

- A. Subscribe to Amazon RDS events to be notified when changes occur with a DB instance, DB snapshot, DB parameter group, or DB security group.
- B. Use the free Amazon CloudWatch service to monitor the performance and health of a DB instance.
- C. All of the items listed will track the performance and health of a database.
- D. View, download, or watch database log files using the Amazon RDS console or Amazon RDS API
- E. You can also query some database log files that are loaded into database tables.

Answer: C

Explanation:

Amazon Relational Database Service (Amazon RDS) is a web service that makes it easier to set up, operate, and scale a relational database in the cloud. It provides cost-efficient, resizable capacity for an industry-standard relational database and manages common database administration tasks.

There are several ways you can track the performance and health of a database or a DB instance. You can:

Use the free Amazon CloudWatch service to monitor the performance and health of a DB instance. Subscribe to Amazon RDS events to be notified when changes occur with a DB instance, DB snapshot, DB parameter group, or DB security group.

View, download, or watch database log files using the Amazon RDS console or Amazon RDS APIs. You can also query some database log files that are loaded into database tables.

Use the AWS CloudTrail service to record AWS calls made by your AWS account. The calls are recorded in log files and stored in an Amazon S3 bucket.

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

NEW QUESTION 42

A user has attached 1 EBS volume to a VPC instance. The user wants to achieve the best fault tolerance of data possible. Which of the below mentioned options can help achieve fault tolerance?

- A. Attach one more volume with RAID 1 configuration.
- B. Attach one more volume with RAID 0 configuration.
- C. Connect multiple volumes and stripe them with RAID 6 configuration.
- D. Use the EBS volume as a root device

Answer: A

Explanation:

The user can join multiple provisioned IOPS volumes together in a RAID 1 configuration to achieve better fault tolerance. RAID 1 does not provide a write performance improvement; it requires more bandwidth than non-RAID configurations since the data is written simultaneously to multiple volumes.

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

NEW QUESTION 47

A user has created a subnet in VPC and launched an EC2 instance within it. The user has not selected the option to assign the IP address while launching the instance. The user has 3 elastic IPs and is trying to assign one of the Elastic IPs to the VPC instance from the console. The console does not show any instance in the IP assignment screen. What is a possible reason that the instance is unavailable in the assigned IP console?

- A. The IP address may be attached to one of the instances
- B. The IP address belongs to a different zone than the subnet zone
- C. The user has not created an internet gateway
- D. The IP addresses belong to EC2 Classic; so they cannot be assigned to VPC

Answer: D

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. When the user is launching an instance he needs to select an option which attaches a public IP to the instance. If the user has not selected the option to attach the public IP then it will only have a private IP when launched. If the user wants to connect to an instance from the internet he should create an elastic IP with VPC. If the elastic IP is a part of EC2 Classic it cannot be assigned to a VPC instance.

Reference: <http://docs.aws.amazon.com/AmazonVPC/latest/GettingStartedGuide/LaunchInstance.html>

NEW QUESTION 51

A user is observing the EC2 CPU utilization metric on CloudWatch. The user has observed some interesting patterns while filtering over the 1 week period for a particular hour. The user wants to zoom that data point to a more granular period. How can the user do that easily with CloudWatch?

- A. The user can zoom a particular period by selecting that period with the mouse and then releasing the mouse
- B. The user can zoom a particular period by specifying the aggregation data for that period
- C. The user can zoom a particular period by double clicking on that period with the mouse
- D. The user can zoom a particular period by specifying the period in the Time Range

Answer: A

Explanation:

Amazon CloudWatch provides the functionality to graph the metric data generated either by the AWS services or the custom metric to make it easier for the user to analyse. The AWS CloudWatch console provides the option to change the granularity of a graph and zoom in to see data over a shorter time period. To zoom, the user has to click in the graph details pane, drag on the graph area for selection, and then release the mouse button.

Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/zoom_in_on_graph.html

NEW QUESTION 53

You need to create an Amazon Machine Image (AMI) for a customer for an application which does not appear to be part of the standard AWS AMI template that you can see in the AWS console. What are the alternative possibilities for creating an AMI on AWS?

- A. You can purchase AMIs from a third party but cannot create your own AMI.
- B. You can purchase AMIs from a third party or can create your own AMI.
- C. Only AWS can create AMIs and you need to wait till it becomes available.
- D. Only AWS can create AMIs and you need to request them to create one for you

Answer: B

Explanation:

You can purchase AMIs from a third party, including AMIs that come with service contracts from organizations such as Red Hat. You can also create an AMI and sell it to other Amazon EC2 users. After you create an AMI, you can keep it private so that only you can use it, or you can share it with a specified list of AWS accounts. You can also make your custom AMI public so that the community can use it. Building a safe, secure, usable AMI for public consumption is a fairly straightforward process, if you follow a few simple guidelines.

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

NEW QUESTION 55

Regarding Amazon Route 53, if your application is running on Amazon EC2 instances in two or more Amazon EC2 regions and if you have more than one Amazon EC2 instance in one or more regions, you can use to route traffic to the correct region and then use to route traffic to instances within the region, based on probabilities that you specify.

- A. weighted-based routing; alias resource record sets
- B. latency-based routing; weighted resource record sets
- C. weighted-based routing; weighted resource record sets
- D. latency-based routing; alias resource record sets

Answer: B

Explanation:

Regarding Amazon Route 53, if your application is running on Amazon EC2 instances in two or more Amazon EC2 regions, and if you have more than one Amazon EC2 instance in one or more regions, you can use latency-based routing to route traffic to the correct region and then use weighted resource record sets to route traffic to instances within the region based on weights that you specify.

Reference: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/Tutorials.html>

NEW QUESTION 57

What is the network performance offered by the c4.8xlarge instance in Amazon EC2?

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

Answer: B

Explanation:

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

NEW QUESTION 60

A user has created a CloudFormation stack. The stack creates AWS services, such as EC2 instances, ELB, AutoScaling, and RDS. While creating the stack it created EC2, ELB and AutoScaling but failed to create RDS. What will CloudFormation do in this scenario?

- A. Rollback all the changes and terminate all the created services
- B. It will wait for the user's input about the error and correct the mistake after the input
- C. CloudFormation can never throw an error after launching a few services since it verifies all the steps before launching
- D. It will warn the user about the error and ask the user to manually create RDS

Answer: A

Explanation:

AWS CloudFormation is an application management tool which provides application modeling, deployment, configuration, management and related activities. The AWS CloudFormation stack is a collection of AWS resources which are created and managed as a single unit when AWS CloudFormation instantiates a template. If any of the services fails to launch, CloudFormation will rollback all the changes and terminate or delete all the created services.

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

NEW QUESTION 63

A major client who has been spending a lot of money on his internet service provider asks you to set up an AWS Direct Connection to try and save him some money. You know he needs high-speed connectivity. Which connection port speeds are available on AWS Direct Connect?

- A. 500Mbps and 1Gbps
- B. 1Gbps and 10Gbps
- C. 100Mbps and 1Gbps
- D. 1Gbps

Answer: B

Explanation:

AWS Direct Connect is a network service that provides an alternative to using the internet to utilize AWS cloud services.

Using AWS Direct Connect, data that would have previously been transported over the Internet can now be delivered through a private network connection between AWS and your datacenter or corporate network.

1Gbps and 10Gbps ports are available. Speeds of 50Mbps, 100Mbps, 200Mbps, 300Mbps, 400Mbps, and 500Mbps can be ordered from any APN partners supporting AWS Direct Connect.

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

NEW QUESTION 64

In Amazon EC2, what is the limit of Reserved Instances per Availability Zone each month?

- A. 5
- B. 20
- C. 50
- D. 10

Answer: B

Explanation:

There are 20 Reserved Instances per Availability Zone in each month.

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

NEW QUESTION 68

The AWS CloudHSM service defines a resource known as a high-availability (HA) , which is a virtual partition that represents a group of partitions, typically distributed between several physical HSMs for high-availability.

- A. proxy group
- B. partition group
- C. functional group
- D. relational group

Answer: B

Explanation:

The AWS CloudHSM service defines a resource known as a high-availability (HA) partition group, which is a virtual partition that represents a group of partitions, typically distributed between several physical HSMs for high-availability.

Reference: <http://docs.aws.amazon.com/cloudhsm/latest/userguide/configuring-ha.html>

NEW QUESTION 72

After moving an E-Commerce website for a client from a dedicated server to AWS you have also set up auto scaling to perform health checks on the instances in your group and replace instances that fail these checks. Your client has come to you with his own health check system that he wants you to use as it has proved to be very useful prior to his site running on AWS. What do you think would be an appropriate response to this given all that you know about auto scaling?

- A. It is not possible to implement your own health check system
- B. You need to use AWS's health check system.
- C. It is not possible to implement your own health check system due to compatibility issues.
- D. It is possible to implement your own health check system and then send the instance's health information directly from your system to Cloud Watch.
- E. It is possible to implement your own health check system and then send the instance's health information directly from your system to Cloud Watch but only in the US East (
- F. Virginia) region.

Answer: C

Explanation:

Auto Scaling periodically performs health checks on the instances in your group and replaces instances that fail these checks. By default, these health checks use the results of EC2 instance status checks to determine the health of an instance. If you use a load balancer with your Auto Scaling group, you can optionally choose to include the results of Elastic Load Balancing health checks.

Auto Scaling marks an instance unhealthy if the calls to the Amazon EC2 action DescribeInstanceStatus returns any other state other than running, the system status shows impaired, or the calls to Elastic Load Balancing action DescribeInstanceHealth returns OutOfService in the instance state field.

After an instance is marked unhealthy because of an Amazon EC2 or Elastic Load Balancing health check, it is scheduled for replacement.

You can customize the health check conducted by your Auto Scaling group by specifying additional checks or by having your own health check system and then sending the instance's health information directly from your system to Auto Scaling.

Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/healthcheck.html>

NEW QUESTION 77

When does the billing of an Amazon EC2 system begin?

- A. It starts when the Status column for your distribution changes from Creating to Deployed.
- B. It starts as soon as you click the create instance option on the main EC2 console.
- C. It starts when your instance reaches 720 instance hours.
- D. It starts when Amazon EC2 initiates the boot sequence of an AM instance

Answer: D

Explanation:

Billing commences when Amazon EC2 initiates the boot sequence of an AM instance. Billing ends when the instance terminates, which could occur through a web services command, by running "shutdown -h", or through instance failure. When you stop an instance, Amazon shuts it down but doesn't charge hourly usage for a stopped instance, or data transfer fees, but charges for the storage for any Amazon EBS volumes.

Reference: <http://aws.amazon.com/ec2/faqs/>

NEW QUESTION 78

Having just set up your first Amazon Virtual Private Cloud (Amazon VPC) network, which defined a default network interface, you decide that you need to create and attach an additional network interface, known as an elastic network interface (ENI) to one of your instances. Which of the following statements is true regarding attaching network interfaces to your instances in your VPC?

- A. You can attach 5 ENIs per instance type.
- B. You can attach as many ENIs as you want.
- C. The number of ENIs you can attach varies by instance type.
- D. You can attach 100 ENIs total regardless of instance type

Answer: C

Explanation:

Each instance in your VPC has a default network interface that is assigned a private IP address from the IP address range of your VPC. You can create and attach an additional network interface, known as an elastic network interface (ENI), to any instance in your VPC. The number of ENIs you can attach varies by instance type.

NEW QUESTION 80

Amazon Elastic Load Balancing is used to manage traffic on a fleet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits.

Which of the following is not an advantage of ELB over an on-premise load balancer?

- A. ELB uses a four-tier, key-based architecture for encryption.
- B. ELB offers clients a single point of contact, and can also serve as the first line of defense against attacks on your network.
- C. ELB takes over the encryption and decryption work from the Amazon EC2 instances and manages it centrally on the load balancer.
- D. ELB supports end-to-end traffic encryption using TLS (previously SSL) on those networks that use secure HTTP (HTTPS) connections.

Answer: A

Explanation:

Amazon Elastic Load Balancing is used to manage traffic on a fleet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits:

Takes over the encryption and decryption work from the Amazon EC2 instances and manages it centrally on the load balancer

Offers clients a single point of contact, and can also serve as the first line of defense against attacks on your network

When used in an Amazon VPC, supports creation and management of security groups associated with your Elastic Load Balancing to provide additional networking and security options

Supports end-to-end traffic encryption using TLS (previously SSL) on those networks that use secure HTTP (HTTPS) connections. When TLS is used, the TLS server certificate used to terminate client connections can be managed centrally on the load balancer, rather than on every individual instance. Reference:

<http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf>

NEW QUESTION 81

You have been asked to build AWS infrastructure for disaster recovery for your local applications and within that you should use an AWS Storage Gateway as part of the solution. Which of the following best describes the function of an AWS Storage Gateway?

- A. Accelerates transferring large amounts of data between the AWS cloud and portable storage devices .
- B. A web service that speeds up distribution of your static and dynamic web content.
- C. Connects an on-premises software appliance with cloud-based storage to provide seamless and secure integration between your on-premises IT environment and AWS's storage infrastructure.
- D. Is a storage service optimized for infrequently used data, or "cold data."

Answer: C

Explanation:

AWS Storage Gateway connects an on-premises software appliance with cloud-based storage to provide seamless integration with data security features between your on-premises IT environment and the Amazon Web Services (AWS) storage infrastructure. You can use the service to store data in the AWS cloud for scalable and cost-effective storage that helps maintain data security. AWS Storage Gateway offers both volume-based and tape-based storage solutions:

Volume gateways Gateway-cached volumes Gateway-stored volumes

Gateway-virtual tape library (VTL)

Reference:

http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_disasterrecovery_07.pdf

NEW QUESTION 83

What happens to data on an ephemeral volume of an EBS-backed EC2 instance if it is terminated or if it fails?

- A. Data is automatically copied to another volume.
- B. The volume snapshot is saved in S3.
- C. Data persists.
- D. Data is delete

Answer: D

Explanation:

Any data on the instance store volumes persists as long as the instance is running, but this data is deleted when the instance is terminated or if it fails (such as if an underlying drive has issues). After an instance store-backed instance fails or terminates, it cannot be restored.

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

NEW QUESTION 86

AWS Identity and Access Management is a web service that enables Amazon Web Services (AWS) customers to manage users and user permissions in AWS. In addition to supporting IAM user policies, some services support resource-based permissions. Which of the following services are supported by resource-based permissions?

- A. Amazon SNS, and Amazon SQS and AWS Direct Connect.
- B. Amazon S3 and Amazon SQS and Amazon ElastiCache.
- C. Amazon S3, Amazon SNS, Amazon SQS, Amazon Glacier and Amazon EBS.
- D. Amazon Glacier, Amazon SNS, and Amazon CloudWatch

Answer: C

Explanation:

In addition to supporting IAM user policies, some services support resource-based permissions, which let you attach policies to the service's resources instead of to IAM users or groups. Resource-based permissions are supported by Amazon S3, Amazon SNS, Amazon SQS, Amazon Glacier and Amazon EBS.

Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_SpecificProducts.htm

NEW QUESTION 91

Which one of the below doesn't affect Amazon CloudFront billing?

- A. Distribution Type
- B. Data Transfer Out
- C. Dedicated IP SSL Certificates
- D. Requests

Answer: A

Explanation:

Amazon CloudFront is a web service for content delivery. CloudFront delivers your content using a global network of edge locations and works seamlessly with Amazon S3 which durably stores the original and definitive versions of your files.

Amazon CloudFront billing is mainly affected by Data Transfer Out

Edge Location Traffic Distribution Requests

Dedicated IP SSL Certificates

Reference: <http://calculator.s3.amazonaws.com/index.html>

NEW QUESTION 94

A friend wants you to set up a small BitTorrent storage area for him on Amazon S3. You tell him it is highly unlikely that AWS would allow such a thing in their infrastructure. However you decide to investigate. Which of the following statements best describes using BitTorrent with Amazon S3?

- A. Amazon S3 does not support the BitTorrent protocol because it is used for pirated software.
- B. You can use the BitTorrent protocol but only for objects that are less than 100 GB in size.
- C. You can use the BitTorrent protocol but you need to ask AWS for specific permissions first.
- D. You can use the BitTorrent protocol but only for objects that are less than 5 GB in size.

Answer: D

Explanation:

BitTorrent is an open, peer-to-peer protocol for distributing files. You can use the BitTorrent protocol to retrieve any publicly-accessible object in Amazon S3. Amazon S3 supports the BitTorrent protocol so that developers can save costs when distributing content at high scale. Amazon S3 is useful for simple, reliable storage of any data. The default distribution mechanism for Amazon S3 data is via client/server download. In client/server distribution, the entire object is transferred point-to-point from Amazon S3 to every authorized user who requests that object. While client/server delivery is appropriate for a wide variety of use cases, it is not optimal for everybody. Specifically, the costs of client/server distribution increase linearly as the number of users downloading objects increases. This can make it expensive to distribute popular objects.

BitTorrent addresses this problem by recruiting the very clients that are downloading the object as distributors themselves: Each client downloads some pieces of the object from Amazon S3 and some from other clients, while simultaneously uploading pieces of the same object to other interested "peers." The benefit for publishers is that for large, popular files the amount of data actually supplied by Amazon S3 can be substantially lower than what it would have been serving the same clients via client/server download. Less data transferred means lower costs for the publisher of the object.

Reference: <http://docs.aws.amazon.com/AmazonS3/latest/dev/S3Torrent.html>

NEW QUESTION 96

A user has defined an AutoScaling termination policy to first delete the instance with the nearest billing hour. AutoScaling has launched 3 instances in the US-East-1A region and 2 instances in the US-East-1B region. One of the instances in the US-East-1B region is running nearest to the billing hour. Which instance will AutoScaling terminate first while executing the termination action?

- A. Random Instance from US-East-1A
- B. Instance with the nearest billing hour in US-East-1B
- C. Instance with the nearest billing hour in US-East-1A
- D. Random instance from US-East-1B

Answer: C

Explanation:

Even though the user has configured the termination policy, before AutoScaling selects an instance to terminate, it first identifies the Availability Zone that has more instances than the other Availability Zones used by the group. Within the selected Availability Zone, it identifies the instance that matches the specified termination policy.

Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/us-termination-policy.html>

NEW QUESTION 98

After deploying a new website for a client on AWS, he asks if you can set it up so that if it fails it can be automatically redirected to a backup website that he has stored on a dedicated server elsewhere. You are wondering whether Amazon Route 53 can do this. Which statement below is correct in regards to Amazon Route 53?

- A. Amazon Route 53 can't help detect an outage
- B. You need to use another service.
- C. Amazon Route 53 can help detect an outage of your website and redirect your end users to alternate locations.
- D. Amazon Route 53 can help detect an outage of your website but can't redirect your end users to alternate locations.
- E. Amazon Route 53 can't help detect an outage of your website, but can redirect your end users to alternate locations.

Answer: B

Explanation:

With DNS Failover, Amazon Route 53 can help detect an outage of your website and redirect your end users to alternate locations where your application is operating properly.

Reference:

<http://aws.amazon.com/about-aws/whats-new/2013/02/11/announcing-dns-failover-for-route-53/>

NEW QUESTION 102

While creating a network in the VPC, which of the following is true of a NAT device?

- A. You have to administer the NAT Gateway Service provided by AWS.
- B. You can choose to use any of the three kinds of NAT devices offered by AWS for special purposes.
- C. You can use a NAT device to enable instances in a private subnet to connect to the Internet.
- D. You are recommended to use AWS NAT instances over NAT gateways, as the instances provide better availability and bandwidth.

Answer: C

Explanation:

You can use a NAT device to enable instances in a private subnet to connect to the Internet (for example, for software updates) or other AWS services, but prevent the Internet from initiating connections with the instances. AWS offers two kinds of NAT devices: a NAT gateway or a NAT instance. We recommend NAT gateways, as they provide better availability and bandwidth over NAT instances. The NAT Gateway service is also a managed service that does not require your administration efforts. A NAT instance is launched from a NAT AM. You can choose to use a NAT instance for special purposes.

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

NEW QUESTION 103

Your manager has asked you to set up a public subnet with instances that can send and receive internet traffic, and a private subnet that can't receive traffic directly from the internet, but can initiate traffic to the internet (and receive responses) through a NAT instance in the public subnet. Hence, the following 3 rules need to be allowed:

Inbound SSH traffic.

Web servers in the public subnet to read and write to MS SQL servers in the private subnet
Inbound RDP traffic from the Microsoft Terminal Services gateway in the public private subnet
What are the respective ports that need to be opened for this?

- A. Ports 22,1433,3389
- B. Ports 21,1433,3389
- C. Ports 25,1433,3389
- D. Ports 22,1343,3999

Answer: A

Explanation:

A network access control list (ACL) is an optional layer of security that acts as a firewall for controlling traffic in and out of a subnet. You might set up network ACLs with rules similar to your security groups in order to add an additional layer of security to your VPC.

The following ports are recommended by AWS for a single subnet with instances that can receive and send Internet traffic and a private subnet that can't receive traffic directly from the Internet. However, it can initiate traffic to the Internet (and receive responses) through a NAT instance in the public subnet. Inbound SSH traffic. Port 22

Web servers in the public subnet to read and write to MS SQL servers in the private subnet. Port 1433
Inbound RDP traffic from the Microsoft Terminal Services gateway in the public private subnet. Port 3389

Reference:

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Appendix_NACLs.html#VPC_Appendix_NACLs_Scenario_2

NEW QUESTION 108

Can I change the EC2 security groups after an instance is launched in EC2-Classic?

- A. Yes, you can change security groups after you launch an instance in EC2-Classic.
- B. No, you cannot change security groups after you launch an instance in EC2-Classic.
- C. Yes, you can only when you remove rules from a security group.
- D. Yes, you can only when you add rules to a security group.

Answer: B

Explanation:

After you launch an instance in EC2-Classic, you can't change its security groups. However, you can add rules to or remove rules from a security group, and those changes are automatically applied to all instances that are associated with the security group.

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

NEW QUESTION 113

Can you encrypt EBS volumes?

- A. Yes, you can enable encryption when you create a new EBS volume using the AWS Management Console, API, or CLI.
- B. No, you should use a third-party software to perform raw block-level encryption of an EBS volume.
- C. Yes, but you must use a third-party API for encrypting data before it's loaded on EBS.
- D. Yes, you can encrypt with the special "ebs_encrypt" command through Amazon API

Answer: A

Explanation:

With Amazon EBS encryption, you can now create an encrypted EBS volume and attach it to a supported instance type. Data on the volume, disk I/O, and snapshots created from the volume are then all encrypted. The encryption occurs on the servers that host the EC2 instances, providing encryption of data as it moves between EC2 instances and EBS storage. EBS encryption is based on the industry standard AES-256 cryptographic algorithm.

To get started, simply enable encryption when you create a new EBS volume using the AWS Management Console, API, or CLI. Amazon EBS encryption is available for all the latest EC2 instances in all commercially available AWS regions.

Reference:

<https://aws.amazon.com/about-aws/whats-new/2014/05/21/Amazon-EBS-encryption-now-available/>

NEW QUESTION 118

In Amazon EC2, you are billed instance-hours when .

- A. your EC2 instance is in a running state
- B. the instance exits from Amazon S3 console
- C. your instance still exits the EC2 console
- D. EC2 instances stop

Answer: A

Explanation:

You are billed instance-hours as long as your EC2 instance is in a running state. Reference: <http://aws.amazon.com/ec2/faqs/>

NEW QUESTION 123

You have some very sensitive data stored on AWS S3 and want to try every possible alternative to keeping it secure in regards to access control. What are the mechanisms available for access control on AWS S3?

- A. (IAM) policies, Access Control Lists (ACLs), bucket policies, and query string authentication.
- B. (IAM) policies, Access Control Lists (ACLs) and bucket policies.
- C. Access Control Lists (ACLs), bucket policies, and query string authentication
- D. (IAM) policies, Access Control Lists (ACLs), bucket policies, query string authentication and encryption.

Answer: A

Explanation:

Amazon S3 supports several mechanisms that give you flexibility to control who can access your data as well as how, when, and where they can access it.

Amazon S3 provides four different access control mechanisms:

AWS Identity and Access Management (IAM) policies, Access Control Lists (ACLs), bucket policies, and query string authentication.

IAM enables organizations to create and manage multiple users under a single AWS account. With IAM policies, you can grant IAM users fine-grained control to your Amazon S3 bucket or objects. You can use ACLs to selectively add (grant) certain permissions on individual objects.

Amazon S3 bucket policies can be used to add or deny permissions across some or all of the objects within a single bucket.

With Query string authentication, you have the ability to share Amazon S3 objects through URLs that are valid for a specified period of time.

NEW QUESTION 126

You have been asked to set up monitoring of your network and you have decided that Cloudwatch would be the best service to use. Amazon CloudWatch monitors your Amazon Web Services (AWS) resources and the applications you run on AWS in real-time. You can use CloudWatch to collect and track metrics, which are the variables you want to measure for your resources and applications. Which of the following items listed can AWS Cloudwatch monitor?

- A. Log files your applications generate.
- B. All of the items listed on this page.
- C. System-wide visibility into resource utilization, application performance, and operational health.
- D. Custom metrics generated by your applications and services .

Answer: B

Explanation:

Amazon CloudWatch can monitor AWS resources such as Amazon EC2 instances, Amazon DynamoDB tables, and Amazon RDS DB instances, as well as custom metrics generated by your applications and services, and any log files your applications generate. You can use Amazon CloudWatch to gain system-wide visibility into resource utilization, application performance, and operational health. You can use these insights to react and keep your application running smoothly.

Reference: <http://aws.amazon.com/cloudwatch/>

NEW QUESTION 127

How can you apply more than 100 rules to an Amazon EC2-Classic?

- A. By adding more security groups
- B. You need to create a default security group specifying your required rules if you need to use more than 100 rules per security group.
- C. By default the Amazon EC2 security groups support 500 rules.
- D. You can't add more than 100 rules to security groups for an Amazon EC2 instance

Answer: D

Explanation:

In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 100 rules to a security group.

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

NEW QUESTION 132

You need to quickly set up an email-sending service because a client needs to start using it in the next hour. Amazon Simple Email Service (Amazon SES) seems to be the logical choice but there are several options available to set it up. Which of the following options to set up SES would best meet the needs of the client?

- A. Amazon SES console
- B. AWS CloudFormation
- C. SMTP Interface
- D. AWS Elastic Beanstalk

Answer: A

Explanation:

Amazon SES is an outbound-only email-sending service that provides an easy, cost-effective way for you to send email.

There are several ways that you can send an email by using Amazon SES. You can use the Amazon SES console, the Simple Mail Transfer Protocol (SMTP) interface, or you can call the Amazon SES API. Amazon SES console—This method is the quickest way to set up your system

Reference: <http://docs.aws.amazon.com/ses/latest/DeveloperGuide/welcome.html>

NEW QUESTION 135

When controlling access to Amazon EC2 resources, each Amazon EBS Snapshot has a attribute that controls which AWS accounts can use the snapshot.

- A. createVolumePermission
- B. LaunchPermission
- C. SharePermission

D. RequestPermission

Answer: A

Explanation:

Each Amazon EBS Snapshot has a createVolumePermission attribute that you can set to one or more AWS Account IDs to share the AM with those AWS Accounts. To allow several AWS Accounts to use a particular EBS snapshot, you can use the snapshots's createVolumePermission attribute to include a list of the accounts that can use it.

Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/UsingIAM.html>

NEW QUESTION 137

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 throughpu

Answer: B

Explanation:

Amazon Simple Queue Service (Amazon SQS) offers a reliable, highly scalable hosted queue for storing messages as they travel between computers. By using Amazon SQS, developers can simply move data between distributed application components performing different tasks, without losing messages or requiring each component to be always available. Amazon SQS makes it easy to build a distributed, decoupled application, working in close conjunction with the Amazon Elastic Compute Cloud (Amazon EC2) and the other AWS infrastructure web services.

What can I do with Amazon SQS?

Amazon SQS is a web service that gives you access to a message queue that can be used to store messages while waiting for a computer to process them. This allows you to quickly build message queuing applications that can be run on any computer on the internet. Since Amazon SQS is highly scalable and you only pay for what you use, you can start small and grow your application as you wish, with no compromise on performance or reliability. This lets you focus on building sophisticated message-based applications, without worrying about how the messages are stored and managed.

You can use Amazon SQS with software applications in various ways. For example, you can: Integrate Amazon SQS with other AWS infrastructure web services to make applications more reliable and filexible.

Use Amazon SQS to create a queue of work where each message is a task that needs to be completed by a process. One or many computers can read tasks from the queue and perform them. Build a microservices architecture, using queues to connect your microservices.

Keep notifications of significant events in a business process in an Amazon SQS queue. Each event can have a corresponding message in a queue, and applications that need to be aware of the event can read and process the messages.

NEW QUESTION 140

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

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

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

The pilot is considered a success and your CEO has managed to get the attention of some potential investors. The business plan requires a deployment of at least 1000 sensors which needs to be supported by the backend. You also need to store sensor data for at least two years to be able to compare year over year Improvements.

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

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

Answer: C

NEW QUESTION 143

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

Answer: A

NEW QUESTION 148

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

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

Answer: D

Explanation:

Amazon RDS Multi-AZ Deployments

Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure (for example, instance hardware failure, storage failure, or network disruption), Amazon RDS performs an automatic failover to the standby, so that you can resume database operations as soon as the failover is complete. Since the endpoint for your DB Instance remains the same after a failover, your application can resume database operation without the need for manual administrative intervention.

Enhanced Durability

Multi-AZ deployments for the MySQL, Oracle, and PostgreSQL engines utilize synchronous physical replication to keep data on the standby up-to-date with the primary. Multi-AZ deployments for the SQL Server engine use synchronous logical replication to achieve the same result, employing SQL Server-native Mirroring technology. Both approaches safeguard your data in the event of a DB Instance failure or loss of an Availability Zone.

If a storage volume on your primary fails in a Multi-AZ deployment, Amazon RDS automatically initiates a failover to the up-to-date standby. Compare this to a Single-AZ deployment: in case of a Single-AZ database failure, a user-initiated point-in-time-restore operation will be required. This operation can take several hours to complete, and any data updates that occurred after the latest restorable time (typically within the last five minutes) will not be available.

Amazon Aurora employs a highly durable, SSD-backed virtualized storage layer purpose-built for database workloads. Amazon Aurora automatically replicates your volume six ways, across three Availability Zones. Amazon Aurora storage is fault-tolerant, transparently handling the loss of up to two copies of data without affecting database write availability and up to three copies without affecting read availability. Amazon Aurora storage is also self-healing. Data blocks and disks are continuously scanned for errors and replaced automatically.

Increased Availability

You also benefit from enhanced database availability when running Multi-AZ deployments. If an Availability Zone failure or DB Instance failure occurs, your availability impact is limited to the time automatic failover takes to complete: typically under one minute for Amazon Aurora and one to two minutes for other database engines (see the RDS FAQ for details).

The availability benefits of Multi-AZ deployments also extend to planned maintenance and backups. In the case of system upgrades like OS patching or DB Instance scaling, these operations are applied first on

the standby, prior to the automatic failover. As a result, your availability impact is, again, only the time required for automatic failover to complete.

Unlike Single-AZ deployments, I/O activity is not suspended on your primary during backup for Multi-AZ deployments for the MySQL, Oracle, and PostgreSQL engines, because the backup is taken from the standby. However, note that you may still experience elevated latencies for a few minutes during backups for Multi-AZ deployments.

On instance failure in Amazon Aurora deployments, Amazon RDS uses RDS Multi-AZ technology to automate failover to one of up to 15 Amazon Aurora Replicas you have created in any of three Availability Zones. If no Amazon Aurora Replicas have been provisioned, in the case of a failure, Amazon RDS will attempt to create a new Amazon Aurora DB instance for you automatically.

NEW QUESTION 150

Your company currently has a 2-tier web application running in an on-premises data center. You have experienced several infrastructure failures in the past two months resulting in significant financial losses. Your CIO is strongly agreeing to move the application to AWS. While working on achieving buy-in from the other company executives, he asks you to develop a disaster recovery plan to help improve Business continuity in the short term. He specifies a target Recovery Time Objective (RTO) of 4 hours and a Recovery Point Objective (RPO) of 1 hour or less. He also asks you to implement the solution within 2 weeks. Your database is 200GB in size and you have a 20Mbps Internet connection.

How would you do this while minimizing costs?

- A. Create an EBS backed private AMI which includes a fresh install of your application
- B. Develop a CloudFormation template which includes your AMI and the required EC2, AutoScaling, and ELB resources to support deploying the application across Multiple- Availability-Zone
- C. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
- D. Deploy your application on EC2 instances within an Auto Scaling group across multiple availability zone
- E. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
- F. Create an EBS backed private AMI which includes a fresh install of your application
- G. Setup a script in your data center to backup the local database every 1 hour and to encrypt and copy the resulting file to an S3 bucket using multi-part upload.
- H. Install your application on a compute-optimized EC2 instance capable of supporting the application's average load
- I. Synchronously replicate transactions from your on-premises database to a database instance in AWS across a secure Direct Connect connection.

Answer: A

Explanation:

Overview of Creating Amazon EBS-Backed AMIs

First, launch an instance from an AMI that's similar to the AMI that you'd like to create. You can connect to your instance and customize it. When the instance is configured correctly, ensure data integrity by

stopping the instance before you create an AMI, then create the image. When you create an Amazon EBS-backed AMI, we automatically register it for you.

Amazon EC2 powers down the instance before creating the AMI to ensure that everything on the instance is stopped and in a consistent state during the creation process. If you're confident that your instance is in a consistent state appropriate for AMI creation, you can tell Amazon EC2 not to power down and reboot the instance. Some file systems, such as XFS, can freeze and unfreeze activity, making it safe to create the image without rebooting the instance.

During the AMI-creation process, Amazon EC2 creates snapshots of your instance's root volume and any other EBS volumes attached to your instance. If any volumes attached to the instance are encrypted, the new AMI only launches successfully on instances that support Amazon EBS encryption. For more information, see Amazon EBS Encryption.

Depending on the size of the volumes, it can take several minutes for the AMI-creation process to complete (sometimes up to 24 hours). You may find it more efficient to create snapshots of your volumes prior to creating your AMI. This way, only small, incremental snapshots need to be created when the AMI is created,

and the process completes more quickly (the total time for snapshot creation remains the same). For more information, see [Creating an Amazon EBS Snapshot](#). After the process completes, you have a new AMI and snapshot created from the root volume of the instance. When you launch an instance using the new AMI, we create a new EBS volume for its root volume using the snapshot. Both the AMI and the snapshot incur charges to your account until you delete them. For more information, see [Deregistering Your AMI](#).

If you add instance-store volumes or EBS volumes to your instance in addition to the root device volume, the block device mapping for the new AMI contains information for these volumes, and the block device mappings for instances that you launch from the new AMI automatically contain information for these volumes. The instance-store volumes specified in the block device mapping for the new instance are new and don't contain any data from the instance store volumes of the instance you used to create the AMI. The data on EBS volumes persists. For more information, see [Block Device Mapping](#).

NEW QUESTION 154

Your system recently experienced down time during the troubleshooting process. You found that a new administrator mistakenly terminated several production EC2 instances.

Which of the following strategies will help prevent a similar situation in the future? The administrator still must be able to:

- launch, start stop, and terminate development resources.
- launch and start production instances.

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

Answer: B

Explanation:

Working with volumes

When an API action requires a caller to specify multiple resources, you must create a policy statement that allows users to access all required resources. If you need to use a Condition element with one or more of these resources, you must create multiple statements as shown in this example.

The following policy allows users to attach volumes with the tag "volume_user=iam-user-name" to instances with the tag "department=dev", and to detach those volumes from those instances. If you attach this policy to an IAM group, the aws:username policy variable gives each IAM user in the group permission to attach or detach volumes from the instances with a tag named volume_user that has his or her IAM user name as a value.

```
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow", "Action": [ "ec2:AttachVolume",
"ec2:DetachVolume" ],
"Resource": "arn :aws:ec2:us-east-1:123456789012:instance/*", "Condition": {
"StringEquals": { "ec2:ResourceTag/department": "dev" }
}
},
{
"Effect": "Allow", "Action": [ "ec2:AttachVolume", "ec2:DetachVolume" ],
"Resource": "arn:aws:ec2:us-east-1:123456789012:volume/*", "Condition": {
"StringEquals": {
"ec2:ResourceTag/volume_user": "${aws:username}" }
}
}
}]
}
```

Launching instances (RunInstances)

The RunInstances API action launches one or more instances. RunInstances requires an AMI and creates an instance; and users can specify a key pair and security group in the request. Launching into EC2-VPC requires a subnet, and creates a network interface. Launching from an Amazon EBS-backed AMI creates a volume. Therefore, the user must have permission to use these Amazon EC2 resources. The caller can also configure the instance using optional parameters to Run Instances, such as the instance type and a subnet. You can create a policy statement that requires users to specify an optional parameter, or restricts users to particular values for a parameter. The examples in this section demonstrate some of the many possible ways that you can control the configuration of an instance that a user can launch.

Note that by default, users don't have permission to describe, start, stop, or terminate the resulting instances. One way to grant the users permission to manage the resulting instances is to create a specific tag for each instance, and then create a statement that enables them to manage instances with that tag. For more information, see [2: Working with instances](#).

a. AMI

The following policy allows users to launch instances using only the AMIs that have the specified tag, "department=dev", associated with them. The users can't launch instances using other AMIs because the Condition element of the first statement requires that users specify an AMI that has this tag. The users also can't launch into a subnet, as the policy does not grant permissions for the subnet and network interface resources. They can, however, launch into EC2-Classic. The second statement uses a wildcard to enable users to create instance resources, and requires users to specify the key pair project_keypair and the security group sg-1a2b3c4d. Users are still able to launch instances without a key pair.

```
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow",
"Action": "ec2:RunInstances", "Resource": [ "arn:aws:ec2:region::image/ami-*" ],
"Condition": { "StringEquals": {
"ec2:ResourceTag/department": "dev" }
}
},
{
"Effect": "Allow",
"Action": "ec2:RunInstances", "Resource": [ "arn:aws:ec2:region:account:instance/*", "arn:aws:ec2:region:account:volume/*",
"arn:aws:ec2:region:account:key-pair/project_keypair",
"arn :aws :ec2: region: account:security-group/sg-1a 2b3c4d" ]
}
}]
}
```

Alternatively, the following policy allows users to launch instances using only the specified AMIs, ami-9e1670f7 and ami-45cf5c3c. The users can't launch an instance using other AMIs (unless another statement grants the users permission to do so), and the users can't launch an instance into a subnet.

```
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow",
```

```
"Action": "ec2:RunInstances", "Resource": [
"arn:aws:ec2:region::image/ami-9e1670f7", "arn:aws:ec2:region::image/ami-45cf5c3c", "arn:aws:ec2:region:account:instance/*",
"arn:aws:ec2:region:account:volume/*", "arn:aws:ec2:region:account:key-pair/*", "arn:aws:ec2:region:account:security-group/*"
]
}
```

Alternatively, the following policy allows users to launch instances from all AMIs owned by Amazon. The Condition element of the first statement tests whether ec2:Owner is amazon. The users can't launch an instance using other AMIs (unless another statement grants the users permission to do so).

```
The users are able to launch an instance into a subnet. "Version": "2012-10-17",
"Statement": [{
"Effect": "Allow",
"Action": "ec2:RunInstances", "Resource": [ "arn:aws:ec2:region::image/ami-*" ],
"Condition": { "StringEquals": { "ec2:Owner": "amazon" }
}
},
{
"Effect": "Allow",
"Action": "ec2:RunInstances", "Resource": [ "arn:aws:ec2:region:account:instance/*", "arn:aws:ec2:region:account:subnet/*",
"arn:aws:ec2:region:account:volume/*",
"arn:aws:ec2:region:account:network-interface/*", "arn:aws:ec2:region:account:key-pair/*", "arn:aws:ec2:region:account:security-group/*"
]
}
]
```

b. Instance type

The following policy allows users to launch instances using only the t2.micro or t2.small instance type, which you might do to control costs. The users can't launch larger instances because the Condition element of the first statement tests whether ec2:InstanceType is either t2.micro or t2.small.

```
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow",
"Action": "ec2:RunInstances", "Resource": [ "arn:aws:ec2:region:account:instance/*" ],
"Condition": { "StringEquals": {
"ec2:InstanceType": ["t2.micro", "t2.small"]
}
}
},
{
"Effect": "Allow",
"Action": "ec2:RunInstances", "Resource": [ "arn:aws:ec2:region::image/ami-*", "arn:aws:ec2:region:account:subnet/*",
"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/*"
]
}
]
```

Alternatively, you can create a policy that denies users permission to launch any instances except t2.micro and t2.small instance types.

```
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Deny",
"Action": "ec2:RunInstances", "Resource": [ "arn:aws:ec2:region:account:instance/*" ],
"Condition": { "StringNotEquals": {
"ec2:InstanceType": ["t2.micro", "t2.small"]
}
}
},
{
"Effect": "Allow",
"Action": "ec2:RunInstances", "Resource": [ "arn:aws:ec2:region::image/ami-*",
"arn:aws:ec2:region:account:network-interface/*", "arn:aws:ec2:region:account:instance/*", "arn:aws:ec2:region:account:subnet/*",
"arn:aws:ec2:region:account:volume/*", "arn:aws:ec2:region:account:key-pair/*", "arn:aws:ec2:region:account:security-group/*"
]
}
]
```

c. Subnet

The following policy allows users to launch instances using only the specified subnet, subnet-12345678. The group can't launch instances into any other subnet (unless another statement grants the users permission to do so). Users are still able to launch instances into EC2-Classic.

```
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow",
"Action": "ec2:RunInstances", "Resource": [
"arn:aws:ec2:region:account:subnet/subnet-12345678",
"arn:aws:ec2:region:account:network-interface/*", "arn:aws:ec2:region:account:instance/*", "arn:aws:ec2:region:account:volume/*",
"arn:aws:ec2:region::image/ami-*", "arn:aws:ec2:region:account:key-pair/*", "arn:aws:ec2:region:account:security-group/*"
]
}
]
```

Alternatively, you could create a policy that denies users permission to launch an instance into any other subnet. The statement does this by denying permission to create a network interface, except where subnet subnet-12345678 is specified. This denial overrides any other policies that are created to allow launching instances into other subnets. Users are still able to launch instances into EC2-Classic.

```
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Deny",
"Action": "ec2:RunInstances", "Resource": [
"arn:aws:ec2:region:account:network-interface/*" ],

```

```
"Condition": { "ArnNotEquals": {
"ec2:Subnet": "arn :aws:ec2:region:account:subnet/subnet-12345678"
}
},
{
"Effect": "Allow",
"Action": "ec2:RunInstances", "Resource": [ "arn:aws:ec2:region::image/ami-*",
"arn:aws:ec2:region:account:network-interface/*", "arn:aws:ec2:region:account:instance/*", "arn:aws:ec2:region:account:subnet/*",
"arn:aws:ec2:region:account:volume/*", "arn:aws:ec2:region:account:key-pair/*", "arn:aws:ec2:region:account:security-group/*"
]
}
}
```

NEW QUESTION 155

You are implementing AWS Direct Connect. You intend to use AWS public service end points such as Amazon 53, across the AWS Direct Connect link. You want other Internet traffic to use your existing link to an Internet Service Provider.

What is the correct way to configure AWS Direct connect for access to services such as Amazon 53?

- A. Configure a public Interface on your AWS Direct Connect link Configure a static route via your AWS Direct Connect link that points to Amazon 53 Advertise a default route to AWS using BGP.
- B. Create a private interface on your AWS Direct Connect link
- C. Configure a static route via your AWS Direct connect link that points to Amazon 53 Configure specific routes to your network in your VPC,
- D. Create a public interface on your AWS Direct Connect link Redistribute BGP routes into your existing routing infrastructure advertise specific routes for your network to AWS.
- E. Create a private interface on your AWS Direct connect link
- F. Redistribute BGP routes into your existing routing infrastructure and advertise a default route to AWS.

Answer: C

NEW QUESTION 157

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

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

Answer: CD

NEW QUESTION 161

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

- A. Use S3 with reduced redundancy to store and serve the scanned files, install the commercial search application on EC2 Instances and configure with auto-scaling and an Elastic Load Balancer.
- B. Model the environment using CloudFormation use an EC2 instance running Apache webserver and an open source search application, stripe multiple standard EB5 volumes together to store the JPEGs and search index.
- C. Use S3 with standard redundancy to store and serve the scanned files, use CloudSearch for queryprocessing, and use Elastic Beanstalk to host the website across multiple availability zones.
- D. Use a single-AZ RD5 My5QL instance to store the search index and the JPEG images use an EC2 instance to serve the website and translate user queries into SQL.
- E. Use a CloudFront download distribution to serve the JPEGs to the end users and Install the current commercial search product, along with a Java Container for the website on EC2 instances and use Route53 with DNS round-robin.

Answer: C

Explanation:

There is no such thing as "Most appropriate" without knowing all your goals. I find your scenarios very fuzzy, since you can obviously mix-n-match between them. I think you should decide by layers instead: Load Balancer Layer: ELB or just DNS, or roll-your-own. (Using DNS+EIPs is slightly cheaper, but less reliable than ELB.)

Storage Layer for 17TB of Images: This is the perfect use case for S3. Off-load all the web requests directly to the relevant JPEGs in S3. Your EC2 boxes just generate links to them.

If your app already serves its own images (not links to images), you might start with EFS. But more than likely, you can just setup a web server to re-write or re-direct all JPEG links to S3 pretty easily.

If you use S3, don't serve directly from the bucket- Serve via a CNAME in domain you control. That way, you can switch in CloudFront easily.

EBS will be way more expensive, and you'll need 2x the drives if you need 2 boxes. Yuck. Consider a smaller storage format. For example, JPEG200 or WebP or other tools might make for smaller images. There is also the DejaVu format from a while back.

Cache Layer: Adding Cloud Front in front of S3 will help people on the other side of the world-- well, possibly. Typical archives follow a power law. The long tail of requests means that most JPEGs won't be requested enough to be in the cache. So you are only speeding up the most popular objects. You can always wait, and switch in CF later after you know your costs better. (In some cases, it can actually lower costs.)

You can also put CloudFront in front of your app, since your archive search results should be fairly static. This will also allow you to run with a smaller instance type, since CF will handle much of the load if you do it right.

Database Layer: A few options:

Use whatever your current server does for now, and replace with something else down the road. Don't under-estimate this approach, sometimes it's better to start now and optimize later.

Use RDS to run MySQL/ Postgres

I'm not as familiar with Elasticsearch I Cloudsearch, but obviously Cloudsearch will be less maintenance+setup.

App Layer:

When creating the app layer from scratch, consider Cloud Formation and/or OpsWorks. It's extra stuff to learn, but helps down the road.

Java+ Tomcat is right up the alley of ElasticBeanstalk. (Basically EC2 + Autoscale + ELB).

Preventing Abuse: When you put something in a public S3 bucket, people will hot-link it from their web pages. If you want to prevent that, your app on the EC2 box can generate signed links to S3 that expire in a few hours. Now everyone will be forced to go thru the app, and the app can apply rate limiting, etc. Saving money:

If you don't mind having downtime:

run everything in one AZ (both DBs and EC2s). You can always add servers and AZs down the road, as long as it's architected to be stateless. In fact, you should use multiple regions if you want it to be really robust.

use Reduced Redundancy in S3 to save a few hundred bucks per month (Someone will have to "go fix it" every time it breaks, including having an off-line copy to repair S3.)

Buy Reserved Instances on your EC2 boxes to make them cheaper. (Start with the RI market and buy a partially used one to get started.) It's just a coupon saying "if you run this type of box in this AZ, you will save on the per-hour costs." You can get 1/2 to 1/3 off easily.

Rewrite the application to use less memory and CPU -that way you can run on fewer/ smaller boxes. (May or may not be worth the investment.)

If your app will be used very infrequently, you will save a lot of money by using Lambda. I'd be worried that it would be quite slow if you tried to run a Java application on it though ..

We're missing some information like load, latency expectations from search, indexing speed, size of the search index, etc. But with what you've given us, I would go with S3 as the storage for the files (S3 rocks. It is really, really awesome). If you're stuck with the commercial search application, then on EC2 instances with autoscaling and an ELB. If you are allowed an alternative search engine, Elasticsearch is probably your best bet. I'd run it on EC2 instead of the AWS Elasticsearch service, as IMHO it's not ready yet. Don't autoscale Elasticsearch automatically though, it'll cause all sorts of issues. I have zero experience with CloudSearch so I can't comment on that. Regardless of which option, I'd use Cloud Formation for all of it.

NEW QUESTION 162

A corporate web application is deployed within an Amazon Virtual Private Cloud (VPC) and is connected to the corporate data center via an IPsec VPN. The application must authenticate against the on-premises LDAP server. After authentication, each logged-in user can only access an Amazon Simple Storage Space (S3) keyspace specific to that user.

Which two approaches can satisfy these objectives? (Choose 2 answers)

- A. Develop an identity broker that authenticates against IAM Security Token service to assume a Lam role in order to get temporary AWS security credentials The application calls the identity broker to get AWS temporary security credentials with access to the appropriate S3 bucket.
- B. The application authenticates against LDAP and retrieves the name of an IAM role associated with the use
- C. The application then calls the IAM Security Token Service to assume that IAM role The application can use the temporary credentials to access the appropriate S3 bucket.
- D. Develop an identity broker that authenticates against LDAP and then calls IAM Security Token Service to get IAM federated user credentials The application calls the identity broker to get IAM federated user credentials with access to the appropriate S3 bucket.
- E. The application authenticates against LDAP the application then calls the AWS identity and Access Management (IAM) Security service to log in to IAM using the LDAP credentials the application can use the IAM temporary credentials to access the appropriate S3 bucket.
- F. The application authenticates against IAM Security Token Service using the LDAP credentials the application uses those temporary AWS security credentials to access the appropriate S3 bucket.

Answer: BC

NEW QUESTION 163

Your department creates regular analytics reports from your company's log files All log data is collected in Amazon S3 and processed by daily Amazon Elastic MapReduce (EMR) jobs that generate daily PDF reports and aggregated tables in CSV format for an Amazon Redshift data warehouse.

Your CFO requests that you optimize the cost structure for this system.

Which of the following alternatives will lower costs without compromising average performance of the system or data integrity for the raw data?

- A. Use reduced redundancy storage (RRS) for all data in S3. Use a combination of Spot Instances and Reserved Instances for Amazon EMR job
- B. Use Reserved Instances for Amazon Redshift.
- C. Use reduced redundancy storage (RRS) for PDF and .csv data in S3. Add Spot Instances to EMR job
- D. Use Spot Instances for Amazon Redshift.
- E. Use reduced redundancy storage (RRS) for PDF and .csv data in Amazon S3. Add Spot Instances to Amazon EMR job
- F. Use Reserved Instances for Amazon Redshift.
- G. Use reduced redundancy storage (RRS) for all data in Amazon S3. Add Spot Instances to Amazon EMR job
- H. Use Reserved Instances for Amazon Redshift.

Answer: C

Explanation:

Using Reduced Redundancy Storage

Amazon S3 stores objects according to their storage class. It assigns the storage class to an object when it is written to Amazon S3. You can assign objects a specific storage class (standard or reduced redundancy) only when you write the objects to an Amazon S3 bucket or when you copy objects that are already stored in Amazon S3. Standard is the default storage class. For information about storage classes, see Object Key and Metadata.

In order to reduce storage costs, you can use reduced redundancy storage for noncritical, reproducible data at lower levels of redundancy than Amazon S3 provides with standard storage. The lower level of redundancy results in less durability and availability, but in many cases, the lower costs can make reduced redundancy storage an acceptable storage solution. For example, it can be a cost effective solution for sharing media content that is durably stored elsewhere. It can also make sense if you are storing thumbnails and other resized images that can be easily reproduced from an original image. Reduced redundancy storage is designed to provide 99.99% durability of objects over a given year.

This durability level corresponds to an average annual expected loss of 0.01% of objects. For example, if you store 10,000 objects using the RRS option, you can, on average, expect to incur an annual loss of a single object per year (0.01% of 10,000 objects).

Note

This annual loss represents an expected average and does not guarantee the loss of less than 0.01% of objects in a given year.

Reduced redundancy storage stores objects on multiple devices across multiple facilities, providing 400 times the durability of a typical disk drive, but it does not replicate objects as many times as Amazon S3 standard storage. In addition, reduced redundancy storage is designed to sustain the loss of data in a single facility. If an object in reduced redundancy storage has been lost, Amazon S3 will return a 405 error on requests made to that object. Amazon S3 also offers notifications for reduced redundancy storage object loss: you can configure your bucket so that when Amazon S3 detects the loss of an RRS object, a notification will be sent

through Amazon Simple Notification Service (Amazon SNS). You can then replace the lost object. To enable notifications, you can use the Amazon S3 console to set the Notifications property of your bucket.

NEW QUESTION 168

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 logs nightly into S3 for further processing with EMR. They are concerned because they found suspicious 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 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 sub net.
- 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 170

You currently operate a web application in the AWS US-East region. The application runs on an autoscaled 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 174

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 175

You are designing an S3 solution that requires HTTPS clients to be authenticated by the Web server 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 HTTPS. Place the Web servers in a Route53 Record Set and configure health checks against all Web servers.
- C. Configure ELB with HTTPS listeners, and place the Web servers behind it.
- D. Configure your web servers as the origins for a CloudFront distribution.
- E. Use custom SSL certificates on your CloudFront distribution.

Answer: AB

NEW QUESTION 177

You are designing a hybrid solution between on-premises infrastructure and Amazon VPC. Your server's on-premises will be communicating with your VPC instances. You will be establishing IPsec tunnels over the internet. You will be using VPN gateways and terminating the IPsec tunnels on AWS supported customer gateways.

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

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

F. Data integrity protection across the Internet

Answer: C0EF

NEW QUESTION 178

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

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

Answer: CEF

NEW QUESTION 183

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

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

Answer: D

Explanation:

Use AWS Identity and Access Management (IAM) to control who in your organization has permission to create and manage security groups and network ACLs (NACL). Isolate the responsibilities and roles for better defense. For example, you can give only your network administrators or security ad min the permission to manage the security groups and restrict other roles.

NEW QUESTION 184

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 2GB 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 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 185

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

- A. Use Cloud Front 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 Cloud Front 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 Cloud Front 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. 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
- F. 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 190

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

Answer: B

Explanation:

The throughput of an Amazon Kinesis stream is designed to scale without limits via increasing the number of shards within a stream. However, there are certain limits you should keep in mind while using Amazon Kinesis Streams:

By default, Records of a stream are accessible for up to 24 hours from the time they are added to the stream. You can raise this limit to up to 7 days by enabling extended data retention.

The maximum size of a data blob (the data payload before Base64-encoding) within one record is 1 megabyte (MB).

Each shard can support up to 1000 PUT records per second.

For more information about other API level limits, see Amazon Kinesis Streams Limits.

NEW QUESTION 193

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

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

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

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

Answer: D

NEW QUESTION 198

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

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

Answer: C

NEW QUESTION 200

A web-startup runs its very successful social news application on Amazon EC2 with an Elastic Load Balancer, an Auto-Scaling group of Java/Tomcat application-servers, and DynamoDB as data store. The main web-application best runs on m2 x large instances since it is highly memory- bound Each new deployment requires semi-automated creation and testing of a new AM for the application servers which takes quite a while and is therefore only done once per week.

Recently, a new chat feature has been implemented in nodejs and wails to be integrated in the architecture. First tests show that the new component is CPU bound Because the company has some experience with using Chef, they decided to streamline the deployment process and use AWS Ops Works as an application life cycle tool to simplify management of the application and reduce the deployment cycles.

What configuration in AWS Ops Works is necessary to integrate the new chat module in the most cost-efficient and flexible way?

- A. Create one AWS OpsWorks stack, create one AWS Ops Works layer, create one custom recipe
- B. Create one AWS OpsWorks stack create two AWS Ops Works layers create one custom recipe
- C. Create two AWS OpsWorks stacks create two AWS Ops Works layers create one custom recipe
- D. Create two AWS OpsWorks stacks create two AWS Ops Works layers create two custom recipe

Answer: C

NEW QUESTION 204

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 oaten process this data and used Rabbit MQAn 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 SNS for passing job messages use Cloud Watch alarms to terminate EC2 worker instances when they become idl
- 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 SNS Once data is processed,
- D. Change the storage class of the S3 objects to Reduced Redundancy Storag
- E. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SNS 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 idl
- G. Once data is processed, change the storage class of the S3 object to Glacier.

Answer: D

NEW QUESTION 209

Before I delete an EBS volume, what can I do if I want to recreate the volume later?

- A. Create a copy of the EBS volume (not a snapshot)
- B. Store a snapshot of the volume
- C. Download the content to an EC2 instance
- D. Back up the data in to a physical disk

Answer: B

NEW QUESTION 214

What does RRS stand for when talking about S3?

- A. Redundancy Removal System
- B. Relational Rights Storage
- C. Regional Rights Standard
- D. Reduced Redundancy Storage

Answer: D

NEW QUESTION 217

What does Amazon SWF stand for?

- A. Simple Web Flow
- B. Simple Work Flow
- C. Simple Wireless Forms
- D. Simple Web Form

Answer: B

NEW QUESTION 221

If I write the below command, what does it do? `ec2-run ami-e3a5408a -n 20 -g appserver`

- A. Start twenty instances as members of appserver group.
- B. Creates 20 rules in the security group named appserver
- C. Terminate twenty instances as members of appserver group.
- D. Start 20 security groups

Answer: A

NEW QUESTION 226

While creating an Amazon RDS DB, your first task is to set up a DB _ that controls what IP addresses or EC2 instances have access to your DB Instance.

- A. Security Pool
- B. Secure Zone
- C. Security Token Pool
- D. Security Group

Answer: D

NEW QUESTION 227

What are the initial settings of a user created security group?

- A. Allow all inbound traffic and Allow no outbound traffic
- B. All low no inbound traffic and All low no outbound traffic
- C. All low no inbound traffic and All low all outbound traffic
- D. Allow all inbound traffic and Allow all outbound traffic

Answer: C

NEW QUESTION 231

True or False: When using IAM to control access to your RDS resources, the key names that can be used are case sensitive. For example, `aws:CurrentTime` is NOT equivalent to `AWS:currenttime`.

- A. TRUE
- B. FALSE

Answer: A

NEW QUESTION 232

Amazon RDS automated backups and DB Snapshots are currently supported for only the __ storage engine

- A. InnoDB
- B. MyISAM

Answer: A

NEW QUESTION 235

Fill in the blanks: The base URI for all requests for instance metadata is _ _

- A. <http://254.169.169.254/latest/>
- B. <http://169.169.254.254/latest/>
- C. <http://127.0.0.1/latest/>
- D. <http://169.254.169.254/latest/>

Answer: D

NEW QUESTION 238

What is the maximum key length of a tag?

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

Answer: D

NEW QUESTION 243

What does a "Domain" refer to in Amazon SWF?

- A. A security group in which only tasks inside can communicate with each other
- B. A special type of worker
- C. A collection of related Workflows
- D. The DNS record for the Amazon SWF service

Answer: C

NEW QUESTION 247

Is creating a Read Replica of another Read Replica supported?

- A. Only in certain regions
- B. Only with MSSQL based RDS
- C. Only for Oracle RDS types
- D. No

Answer: D

NEW QUESTION 249

What is the maximum write throughput I can provision for a single Dynamic DB table?

- A. 1,000 write capacity units
- B. 100,000 write capacity units
- C. Dynamic DB is designed to scale without limits, but if you go beyond 10,000 you have to contact AWS first.
- D. 10,000 write capacity units

Answer: C

NEW QUESTION 254

True or False: Manually created DB Snapshots are deleted after the DB Instance is deleted.

- A. TRUE
- B. FALSE

Answer: A

NEW QUESTION 259

What does specifying the mapping `/dev/sdc=none` when launching an instance do?

- A. Prevents `/dev/sdc` from creating the instance.
- B. Prevents `/dev/sdc` from deleting the instance.
- C. Set the value of `/dev/sdc` to 'zero'.
- D. Prevents `/dev/sdc` from attaching to the instance.

Answer: D

NEW QUESTION 264

Is there a limit to how many groups a user can be in?

- A. Yes for all users
- B. Yes for all users except root
- C. No
- D. Yes unless special permission granted

Answer: A

NEW QUESTION 269

While launching an RDS DB instance, on which page I can select the Availability Zone?

- A. REVIEW
- B. DB INSTANCE DETAILS
- C. MANAGEMENT OPTIONS
- D. ADDITIONAL CONFIGURATION

Answer: D

NEW QUESTION 271

Provisioned IOPS Costs: you are charged for the IOPS and storage whether or not you use them in a given month.

- A. FALSE
- B. TRUE

Answer: B

NEW QUESTION 275

IAM's Policy Evaluation Logic always starts with a default _ for every request, except for those that use the AWS account's root security credentials b

- A. Permit
- B. Deny
- C. Cancel

Answer: B

NEW QUESTION 278

By default, when an EBS volume is attached to a Windows instance, it may show up as any drive letter on the instance. You can change the settings of the _ Senrice to set the drive letters of the EBS volumes per your specifications.

- A. EBS Config Senrice
- B. AMI Config Senrice
- C. Ec2 Config Senrice
- D. Ec2-AMI Config Senrice

Answer: C

NEW QUESTION 280

For each DB Instance class, what is the maximum size of associated storage capacity?

- A. 5GB
- B. 1 TB
- C. 2TB
- D. 500GB

Answer: B

NEW QUESTION 281

SQL Sewer _ store log ins and passwords in the master database.

- A. can be configured to but by default does not
- B. doesn't
- C. does

Answer: C

NEW QUESTION 282

What is Oracle SQL Developer?

- A. An AWS developer who is an expert in Amazon RDS using both the Oracle and SQL Server DB engines
- B. A graphical Java tool distributed without cost by Oracle.
- C. It is a variant of the SQL Sewer Management Studio designed by Microsoft to support Oracle DBMS functionalities
- D. A different DBMS released by Microsoft free of cost

Answer: B

NEW QUESTION 286

Does Amazon RDS allow direct host access via Telnet, Secure Shell (SSH), or Windows Remote Desktop Connection?

- A. Yes
- B. No
- C. Depends on if it is in VPC or not

Answer: B

NEW QUESTION 287

Using Amazon IAM, can I give permission based on organizational groups?

- A. Yes but only in certain cases
- B. No
- C. Yes always

Answer: C

NEW QUESTION 291

What is an isolated database environment running in the cloud (Amazon RDS) called?

- A. DB Instance
- B. DB Sewer
- C. DB Unit
- D. DB Volume

Answer: A

NEW QUESTION 296

While signing in REST/ Query requests, for additional security, you should transmit your requests using Secure Sockets Layer (SSL) by using

- A. HTIP
- B. Internet Protocol Security(IPsec)
- C. TLS (Transport Layer Security)
- D. HTIPS

Answer: D

NEW QUESTION 300

When running my DB Instance as a Multi-AZ deployment, can I use the standby for read or write operations?

- A. Yes
- B. Only with MSSQL based RDS
- C. Only for Oracle RDS instances
- D. No

Answer: D

NEW QUESTION 304

Fill in the blanks: "To ensure failover capabilities, consider using a _ for incoming traffic on a network interface".

- A. primary public IP
- B. secondary private I P
- C. secondary public I P
- D. add on secondary IP

Answer: B

NEW QUESTION 306

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AWS-Solution-Architect-Associate Exam with Our Prep Materials Via below:

<https://www.certleader.com/AWS-Solution-Architect-Associate-dumps.html>