

Amazon-Web-Services

Exam Questions SAA-C01

AWS Certified Solutions Architect - Associate



NEW QUESTION 1

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

What AWS architecture would you recommend?

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

Answer: A

Explanation:

In question we have keywords `scalable` and company wants to `move systems` to AWS, which is best suited for Auto-scaling group.

<https://aws.amazon.com/blogs/security/writing-iam-policies-grant-access-to-user-specific-foldersin-an-amazon-s3-bucket/>

NEW QUESTION 2

Your startup wants to implement an order fulfillment process for selling a personalized gadget that needs an average of 3-4 days to produce with some orders taking up to 6 months you expect 10 orders per day on your first day. 1000 orders per day after 6 months and 10,000 orders after 12 months.

Orders coming in are checked for consistency men dispatched to your manufacturing plant for production quality control packaging shipment and payment processing If the product does not meet the quality standards at any stage of the process employees may force the process to repeat a step Customers are notified via email about order status and any critical issues with their orders such as payment failure.

Your case architecture includes AWS Elastic Beanstalk for your website with an RDS MySQL instance for customer data and orders.

How can you implement the order fulfillment process while making sure that the emails are delivered reliably?

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

Answer: C

Explanation:

http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_ecommerce_checkout_13.pdf

NEW QUESTION 3

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

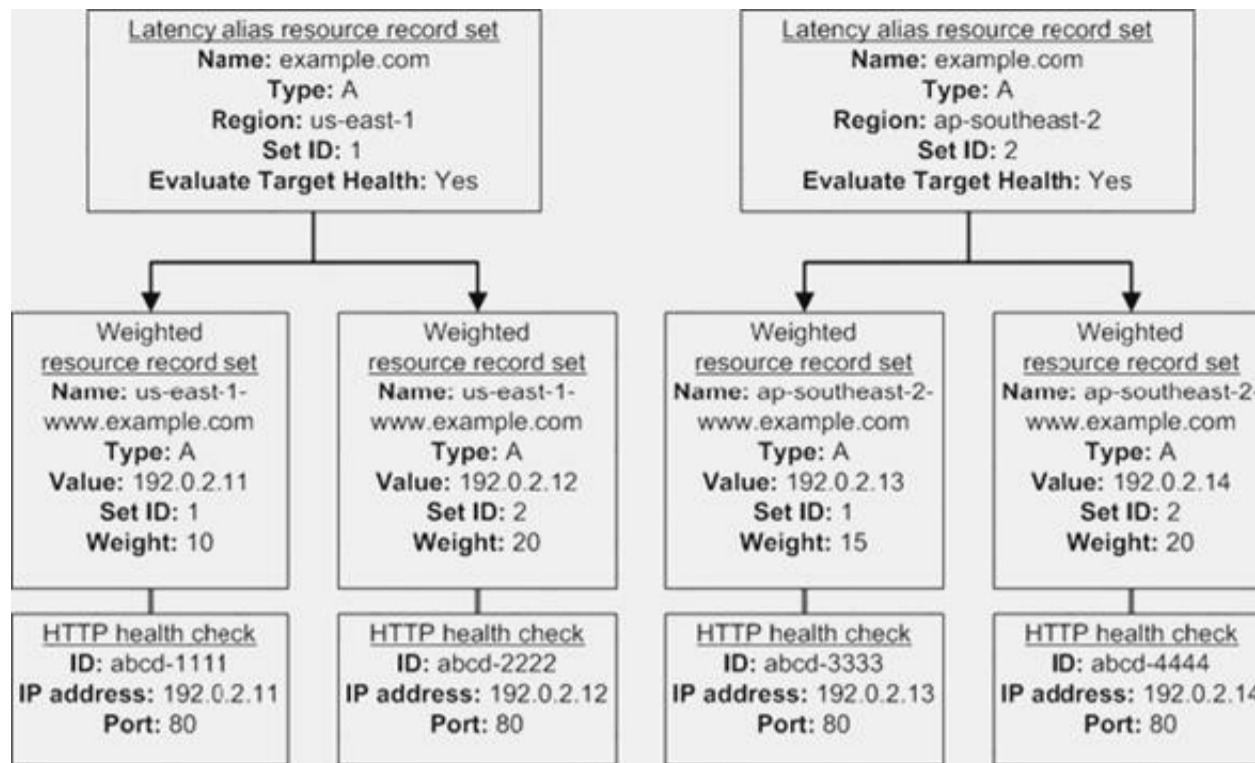
- A. Latency resource record sets cannot be used in combination with weighted resource record sets.
- B. You did not setup an HTTP health check tor one or more of the weighted resource record sets associated with me disabled web servers.
- C. The value of the weight associated with the latency alias resource record set in the region with the disabled servers is higher than the weight for the other region.
- D. One of the two working web servers in the other region did not pass its HTTP health check.
- E. You did not set "Evaluate Target Health" to "Yes" on the latency alias resource record set associated with example com in the region where you disabled the servers.

Answer: BE

Explanation:

How Health Checks Work in Complex Amazon Route 53 Configurations

Checking the health of resources in complex configurations works much the same way as in simple configurations. However, in complex configurations, you use a combination of alias resource record sets (including weighted alias, latency alias, and failover alias) and nonalias resource record sets to build a decision tree that gives you greater control over how Amazon Route 53 responds to requests. For more information, see How Health Checks Work in Simple Amazon Route 53 Configurations. For example, you might use latency alias resource record sets to select a region close to a user and use weighted resource record sets for two or more resources within each region to protect against the failure of a single endpoint or an Availability Zone. The following diagram shows this configuration.



Here's how Amazon EC2 and Amazon Route 53 are configured:

You have Amazon EC2 instances in two regions, us-east-1 and ap-southeast-2. You want Amazon Route 53 to respond to queries by using the resource record sets in the region that provides the lowest latency for your customers, so you create a latency alias resource record set for each region. (You create the latency alias resource record sets after you create resource record sets for the individual Amazon EC2 instances.)

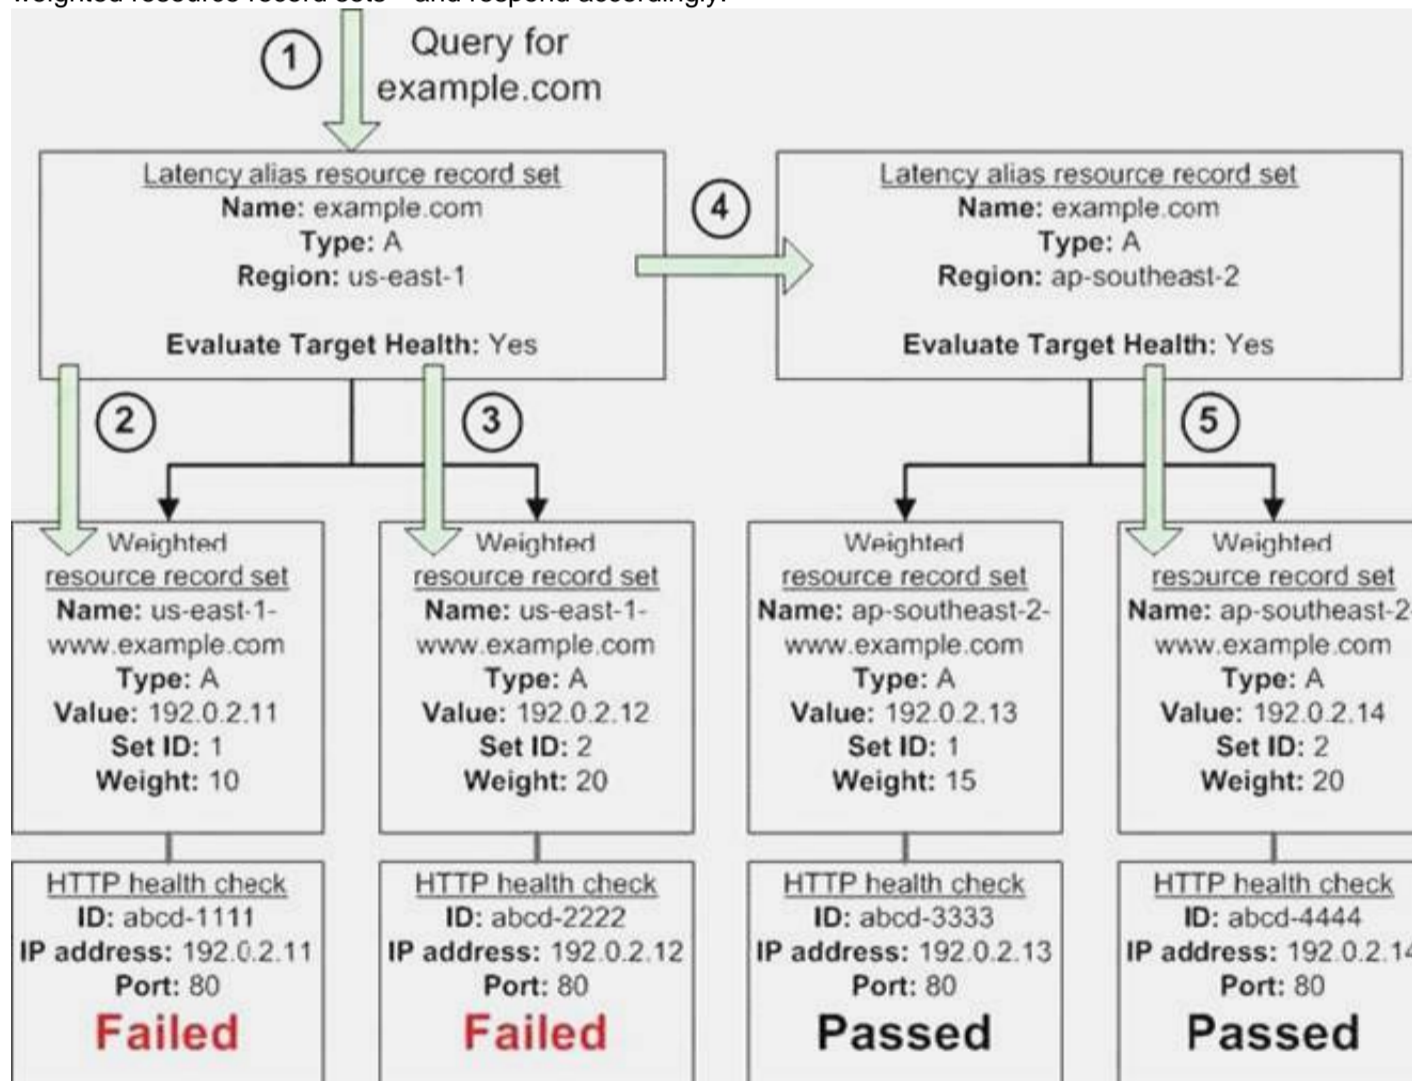
Within each region, you have two Amazon EC2 instances. You create a weighted resource record set for each instance. The name and the type are the same for both of the weighted resource record sets in each region.

When you have multiple resources in a region, you can create weighted or failover resource record sets for your resources. You can also create even more complex configurations by creating weighted alias or failover alias resource record sets that, in turn, refer to multiple resources.

Each weighted resource record set has an associated health check. The IP address for each health check matches the IP address for the corresponding resource record set. This isn't required, but it's the most common configuration.

For both latency alias resource record sets, you set the value of Evaluate Target Health to Yes.

You use the Evaluate Target Health setting for each latency alias resource record set to make Amazon Route 53 evaluate the health of the alias targets—the weighted resource record sets—and respond accordingly.



The preceding diagram illustrates the following sequence of events:

Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region. Amazon Route 53 selects a weighted resource record set based on weight. Evaluate Target Health is Yes for the latency alias resource record set, so Amazon Route 53 checks the health of the selected weighted resource record set.

The health check failed, so Amazon Route 53 chooses another weighted resource record set based on weight and checks its health. That resource record set also is unhealthy.

Amazon Route 53 backs out of that branch of the tree, looks for the latency alias resource record set with the next-best latency, and chooses the resource record set for ap-southeast-2.

Amazon Route 53 again selects a resource record set based on weight, and then checks the health of the selected resource record set. The health check passed, so Amazon Route 53 returns the applicable value in response to the query.

What Happens When You Associate a Health Check with an Alias Resource Record Set?

You can associate a health check with an alias resource record set instead of or in addition to setting the value of Evaluate Target Health to Yes. However, it's generally more useful if Amazon Route 53 responds to queries based on the health of the underlying resources—the HTTP servers, database servers, and other resources that your alias resource record sets refer to. For example, suppose the following configuration:

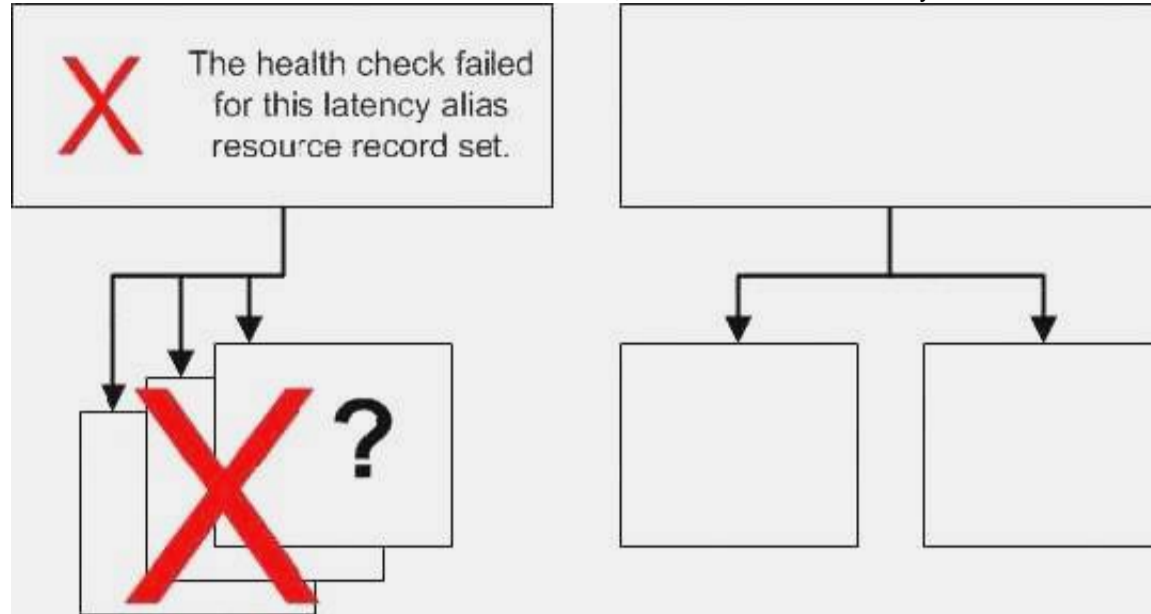
You assign a health check to a latency alias resource record set for which the alias target is a group of weighted resource record sets.

You set the value of Evaluate Target Health to Yes for the latency alias resource record set.

In this configuration, both of the following must be true before Amazon Route 53 will return the applicable value for a weighted resource record set:

The health check associated with the latency alias resource record set must pass.

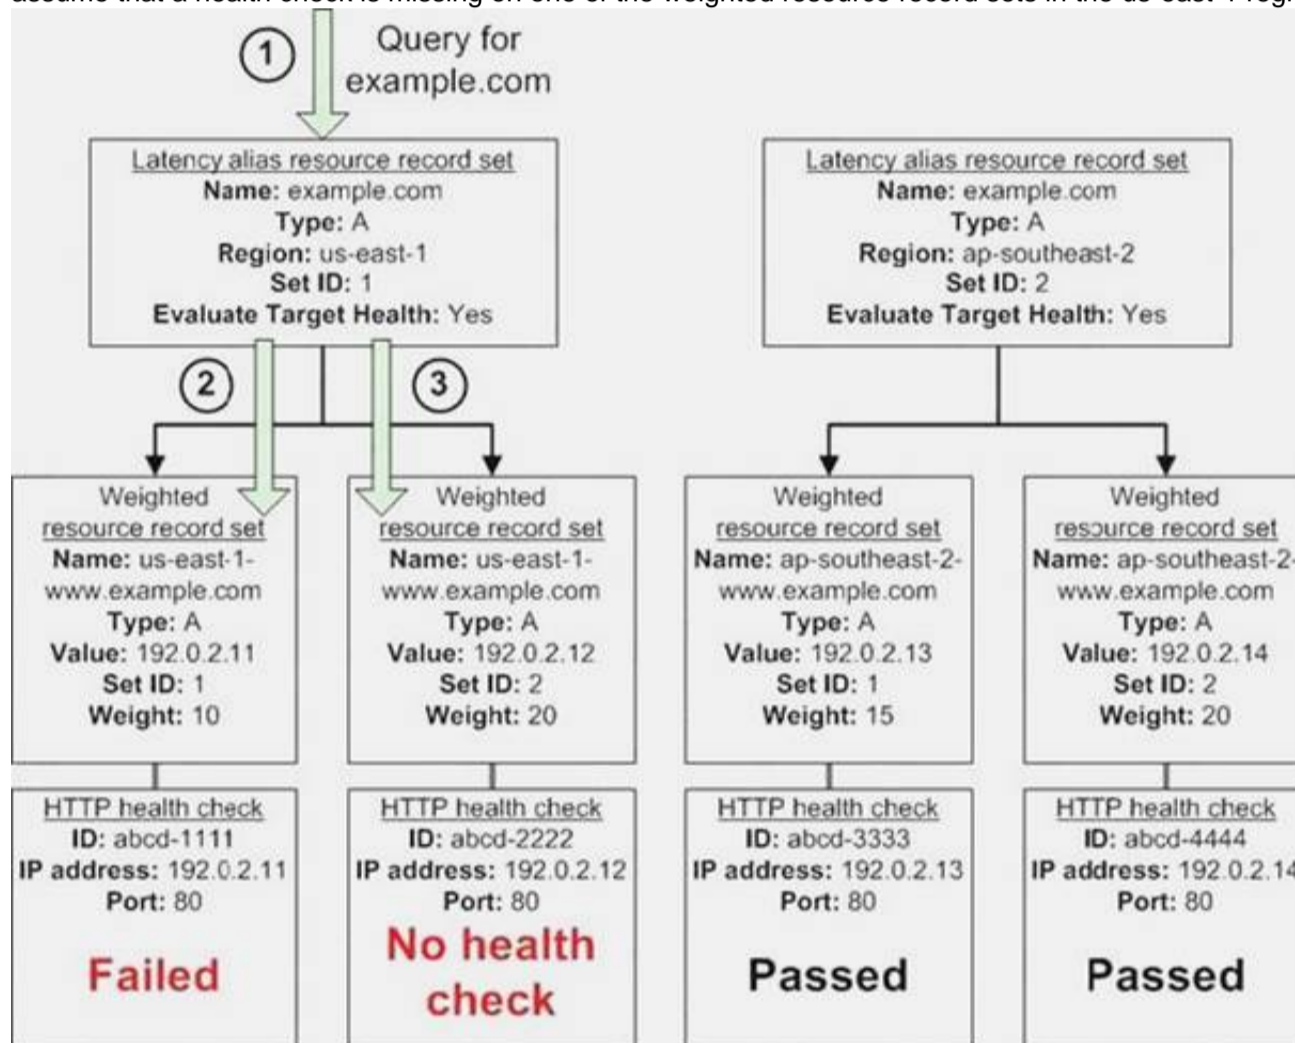
At least one weighted resource record set must be considered healthy, either because it's associated with a health check that passes or because it's not associated with a health check. In the latter case, Amazon Route 53 always considers the weighted resource record set healthy.



If the health check for the latency alias resource record set fails, Amazon Route 53 stops responding to queries using any of the weighted resource record sets in the alias target, even if they're all healthy. Amazon Route 53 doesn't know the status of the weighted resource record sets because it never looks past the failed health check on the alias resource record set.

What Happens When You Omit Health Checks?

In a complex configuration, it's important to associate health checks with all of the non-alias resource record sets. Let's return to the preceding example, but assume that a health check is missing on one of the weighted resource record sets in the us-east-1 region:



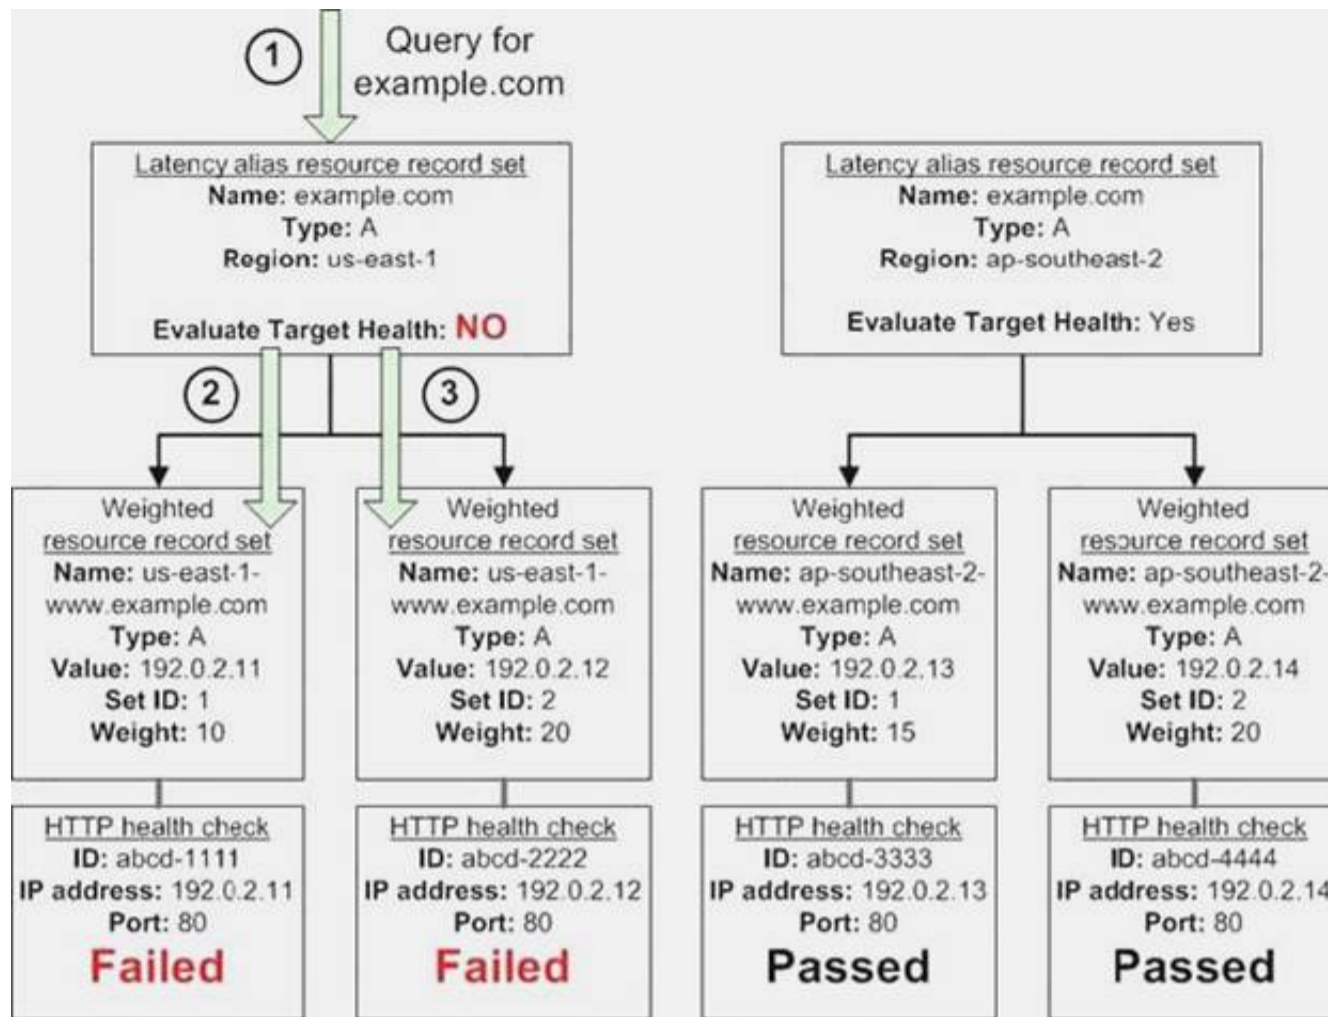
Here's what happens when you omit a health check on a non-alias resource record set in this configuration:

Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region. Amazon Route 53 looks up the alias target for the latency alias resource record set, and checks the status of the corresponding health checks. The health check for one weighted resource record set failed, so that resource record set is omitted from consideration.

The other weighted resource record set in the alias target for the us-east-1 region has no health check. The corresponding resource might or might not be healthy, but without a health check, Amazon Route 53 has no way to know. Amazon Route 53 assumes that the resource is healthy and returns the applicable value in response to the query.

What Happens When You Set Evaluate Target Health to No?

In general, you also want to set Evaluate Target Health to Yes for all of the alias resource record sets. In the following example, all of the weighted resource record sets have associated health checks, but Evaluate Target Health is set to No for the latency alias resource record set for the us-east-1 region:



Here's what happens when you set Evaluate Target Health to No for an alias resource record set in this configuration: Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region. Amazon Route 53 determines what the alias target is for the latency alias resource record set, and checks the corresponding health checks. They're both failing. Because the value of Evaluate Target Health is No for the latency alias resource record set for the useast- 1 region, Amazon Route 53 must choose one resource record set in this branch instead of backing out of the branch and looking for a healthy resource record set in the ap-southeast-2 region.

NEW QUESTION 4

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: D

Explanation:

Reserved Instances (a.k.a. Reserved Nodes) are appropriate for steady-state production workloads, and offer significant discounts over On-Demand pricing.
<https://aws.amazon.com/redshift>

Q: What are some EMR best practices?

If you are running EMR in production you should specify an AMI version, Hive version, Pig version, etc. to make sure the version does not change unexpectedly (e.g. when EMR later adds support for a newer version). If your cluster is mission critical, only use Spot instances for task nodes because if the Spot price increases you may lose the instances. In development, use logging and enable debugging to spot and correct errors faster. If you are using GZIP, keep your file size to 1–2 GB because GZIP files cannot be split. Click here to download the white paper on Amazon EMR best practices. <https://aws.amazon.com/elasticmapreduce/faqs>

NEW QUESTION 5

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 selecte
- B. Use IAM roles S3 bucket policies and Multi Factor Authentication (MFA). Delete on the S3 bucket that stores your logs.
- C. Create a new CloudTrail with one new S3 bucket to store the log
- D. Configure SNS to send log file delivery notifications to your management syste
- E. Use IAM roles and S3 bucket policies on the S3 bucket mat stores your logs.
- F. Create a new CloudTrail trail with an existing S3 bucket to store the logs and with the global services option selecte
- G. Use S3 ACLs and Multi Factor Authentication (MFA). Delete on the S3 bucket that stores your logs.
- H. 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 tool
- I. Use IAM roles and S3 bucket policies on the S3 buckets that store your logs.

Answer: A

NEW QUESTION 6

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

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

Answer: C

NEW QUESTION 7

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 three.)

- 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: CDE

NEW QUESTION 8

You require the ability to analyze a customer's clickstream data on a website so they can do behavioral analysis. Your customer needs to know what sequence of pages and ads their customer clicked on. This data will be used in real time to modify the page layouts as customers click through the site to increase stickiness and advertising click-through. Which option meets the requirements for captioning and analyzing this data?

- A. Log clicks in weblogs by URL store to Amazon S3, and then analyze with Elastic MapReduce
- B. Push web clicks by session to Amazon Kinesis and analyze behavior using Kinesis workers
- C. Write click events directly to Amazon Redshift and then analyze with SQL
- D. Publish web clicks by session to an Amazon SQS queue then periodically drain these events to Amazon RDS and analyze with sol

Answer: B

NEW QUESTION 9

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

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

Answer: B

Explanation:

<https://aws.amazon.com/blogs/aws/fine-grained-access-control-for-amazon-dynamodb/> Here are some of the things that you can build using fine-grained access control:

A mobile app that displays information for nearby airports, based on the user's location. The app can access and display attributes such as airline names, arrival times, and flight numbers. However,

it cannot access or display pilot names or passenger counts.

A mobile game which stores high scores for all users in a single table. Each user can update their own scores, but has no access to the other ones.

NEW QUESTION 10

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

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

Answer: C

NEW QUESTION 10

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

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

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

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

Answer: A

NEW QUESTION 13

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

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

Answer: C

Explanation:

Amazon RDS allows you to use read replicas with Multi-AZ deployments. In Multi-AZ deployments for MySQL, Oracle, SQL Server, and PostgreSQL, the data in your primary DB Instance is synchronously replicated to a standby instance in a different Availability Zone (AZ). Because of their synchronous replication, Multi-AZ deployments for these engines offer greater data durability benefits than do read replicas. (In all Amazon RDS for Aurora deployments, your data is automatically replicated across 3 Availability Zones.)

You can use Multi-AZ deployments and read replicas in conjunction to enjoy the complementary benefits of each. You can simply specify that a given Multi-AZ deployment is the source DB Instance for your Read replicas. That way you gain both the data durability and availability benefits of Multi-AZ deployments and the read scaling benefits of read replicas.

Note that for Multi-AZ deployments, you have the option to create your read replica in an AZ other than that of the primary and the standby for even more redundancy. You can identify the AZ corresponding to your standby by looking at the "Secondary Zone" field of your DB Instance in the AWS Management Console.

NEW QUESTION 14

What does Amazon S3 stand for?

- A. Simple Storage Solution.
- B. Storage Storage Storage (triple redundancy Storage).
- C. Storage Server Solution.
- D. Simple Storage Service

Answer: D

Explanation:

Amazon Simple Storage Service (Amazon S3) is storage for the Internet. It provides a simple interface to manage scalable, reliable, and low latency data storage service over the Internet. <http://docs.aws.amazon.com/AmazonS3/latest/gsg/GetStartedWithS3.html>

NEW QUESTION 19

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

Explanation:

After you no longer need an Amazon EBS volume, you can delete it. After deletion, its data is gone and the volume can't be attached to any instance. However,

before deletion, you can store a snapshot of the volume, which you can use to re-create the volume later.
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-deleting-volume.html>

NEW QUESTION 24

If I want an instance to have a public IP address, which IP address should I use?

- A. Elastic IP Address
- B. Class B IP Address
- C. Class A IP Address
- D. Dynamic IP Address

Answer: A

NEW QUESTION 26

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

Explanation:

In Amazon S3, RRS stands for Reduced Redundancy Storage. 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. <http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingRRS.html>

NEW QUESTION 28

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 29

When you run a DB Instance as a Multi-AZ deployment, the " " serves database writes and reads

- A. secondary
- B. backup
- C. stand by
- D. primary

Answer: D

NEW QUESTION 31

By default, EBS volumes that are created and attached to an instance at launch are deleted when that instance is terminated. You can modify this behavior by changing the value of the flag to false when you launch the instance

- A. DeleteOnTermination
- B. RemoveOnDeletion
- C. RemoveOnTermination
- D. TerminateOnDeletion

Answer: A

Explanation:

By default, Amazon EBS root device volumes are automatically deleted when the instance terminates. However, by default, any additional EBS volumes that you attach at launch, or any EBS volumes that you attach to an existing instance persist even after the instance terminates.

This behavior is controlled by the volume's DeleteOnTermination attribute, which you can modify.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/terminating-instances.html>

NEW QUESTION 36

What will be the status of the snapshot until the snapshot is complete.

- A. running
- B. working
- C. progressing
- D. pending

Answer: D

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html>

Creating an Amazon EBS Snapshot

After writing data to an EBS volume, you can periodically create a snapshot of the volume to use as a baseline for new volumes or for data backup. If you make periodic snapshots of a volume, the snapshots are incremental so that only the blocks on the device that have changed after your last snapshot are saved in the new snapshot. Even though snapshots are saved incrementally, the snapshot deletion process is designed so that you need to retain only the most recent snapshot in order to restore the volume.

Snapshots occur asynchronously; the point-in-time snapshot is created immediately, but the status of the snapshot is **pending** until the snapshot is complete (when all of the modified blocks have been transferred to Amazon S3), which can take several hours for large initial snapshots or subsequent snapshots where many blocks have changed. While it is completing, an in-progress snapshot is not affected by ongoing reads and writes to the volume.

NEW QUESTION 40

Can we attach an EBS volume to more than one EC2 instance at the same time?

- A. No
- B. Yes.
- C. Only EC2-optimized EBS volumes.
- D. Only in read mod

Answer: A

NEW QUESTION 45

Groups can't ____.

- A. be nested more than 3 levels
- B. be nested at all
- C. be nested more than 4 levels
- D. be nested more than 2 levels

Answer: B

Explanation:

Groups can't be nested; they can contain only users, not other groups. http://docs.aws.amazon.com/IAM/latest/UserGuide/id_groups.html

NEW QUESTION 50

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

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

Answer: B

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/raid-config.html> raid 0 and 1 are the common types. Raid 5 and 6 are not recommended because of the extended stripe. If you encounter this question on the exam I suspect the answer options will be different.

Raid 1 Disadvantage

Does not provide a write performance improvement; requires more Amazon EC2 to Amazon EBS bandwidth than non- RAID configurations because the data is written to multiple volumes simultaneously.

Raid 0 Disadvantage

Performance of the stripe is limited to the worst performing volume in the set. Loss of a single volume results in a complete data loss for the array.

Raid 5 and 6 notes

RAID 5 and RAID 6 are not recommended for Amazon EBS because the parity write operations of these RAID modes consume some of the IOPS available to your volumes. Depending on the configuration of your RAID array, these RAID modes provide 20-30% fewer usable IOPS than a RAID 0 configuration. Increased cost is a factor with these RAID modes as well; when using identical volume sizes and speeds, a 2-volume RAID 0 array can outperform a 4-volume RAID 6 array that costs twice as much.

NEW QUESTION 51

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

Explanation:

<https://aws.amazon.com/dynamodb/faqs/>

Q: Is there a **limit** to how much throughput I can get out of a single table?

No, you can increase the throughput you have provisioned for your table using UpdateTable API or in the AWS Management Console. DynamoDB is able to operate at massive scale and there is no theoretical limit on the maximum throughput you can achieve. DynamoDB automatically divides your table across multiple partitions, where each partition is an independent parallel computation unit. DynamoDB can achieve increasingly high throughput rates by adding more partitions.

If you wish to exceed throughput rates of 10,000 writes/second or 10,000 reads/second, you must first contact Amazon through this **online form**.

NEW QUESTION 55

What is the durability of S3 RRS?

- A. 99.99%
- B. 99.95%
- C. 99.995%
- D. 99.999999999%

Answer: A

Explanation:

RRS = Reduced Redundancy Storage

	Standard	Standard - Infrequent Access	Reduced Redundancy Storage
Durability	99.999999999%	99.999999999%	99.99%

NEW QUESTION 58

True or False: When you perform a restore operation to a point in time or from a DB Snapshot, a new DB Instance is created with a new endpoint.

- A. FALSE
- B. TRUE

Answer: B

Explanation:

Restoring From a DB Snapshot

Amazon RDS creates a storage volume snapshot of your DB instance, backing up the entire DB instance and not just individual databases. You can create a DB instance by restoring from this DB snapshot. When you restore the DB instance, you provide the name of the DB snapshot to restore from, and then provide a name for the new DB instance that is created from the restore. You cannot restore from a DB snapshot to an existing DB instance; a new DB instance is created when you restore. http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RestoreFromSnapshot.html

NEW QUESTION 61

How many types of block devices does Amazon EC2 support?

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

Answer: A

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/block-device-mapping-concepts.html> Amazon EC2 supports two types of block devices:

Instance store volumes (virtual devices whose underlying hardware is physically attached to the host computer for the instance)

EBS volumes (remote storage devices)

A block device mapping defines the block devices (instance store volumes and EBS volumes) to attach to an instance.

Block Device Mapping Concepts

A *block device* is a storage device that moves data in sequences of bytes or bits (blocks). These devices support random access and generally use buffered I/O. Examples include hard disks, CD-ROM drives, and flash drives. A block device can be physically attached to a computer or accessed remotely as if it were physically attached to the computer. Amazon EC2 supports **two types** of block devices:

- Instance store volumes (virtual devices whose underlying hardware is physically attached to the host computer for the instance)
- EBS volumes (remote storage devices)

NEW QUESTION 63

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 ____ Service to set the drive letters of the EBS volumes per your specifications.

- A. EBSConfig Service
- B. AMIConfig Service
- C. Ec2Config Service
- D. Ec2-AMIConfig Service

Answer: C

Explanation:

Ec2Config Service is like sysprep and used specifically for windows instances. You can change parameters in OS before launching.

NEW QUESTION 65

SQL Server _____ store logins and passwords in the master database.

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

Answer: C

Explanation:

There are two authentications Windows authentication

The credentials for which are not stored in SQL Server database and managed by windows/AD. There would be entry for windows authenticated logins in master database with respective SID but password would be with Active directory.

SQL Server authentication.

For 2nd we have password stored in hash format you can see it from sys.sql_logins. The information about SQL server logins are stored in master database and each login has SID receptive to it. Only SA login has same SID no matter what server it is. That is why when you move database by backup restore mechanism users are moved not logins and you finally have to create logins(if already not there) and map it to users. This is generally called as troubleshooting orphaned users

NEW QUESTION 67

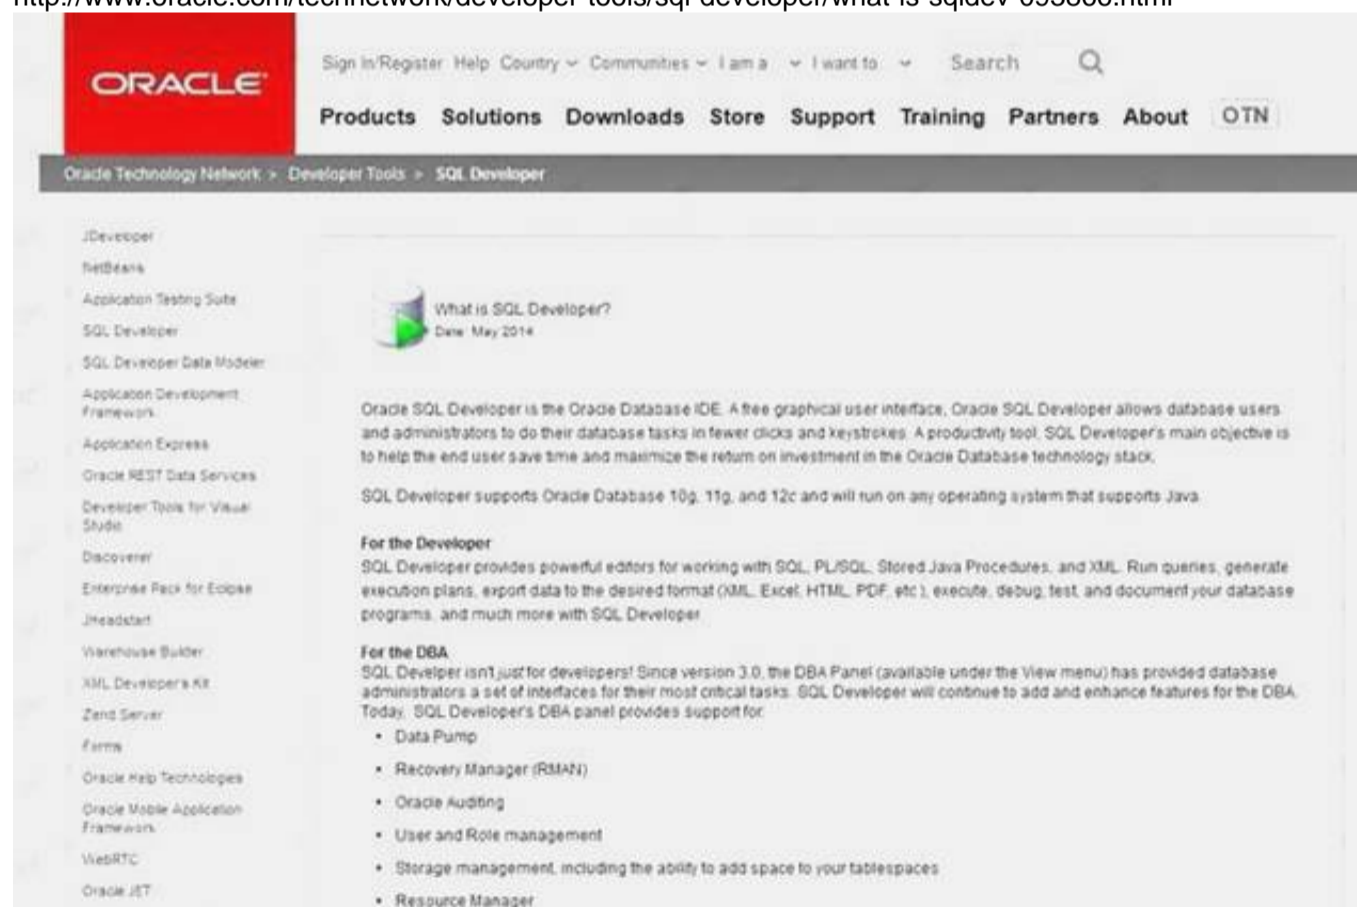
What is Oracle SQL Developer?

- 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 Server Management Studio designed by Microsoft to support Oracle DBMS functionalities
- D. A different DBMS released by Microsoft free of cost

Answer: B

Explanation:

<http://www.oracle.com/technetwork/developer-tools/sql-developer/what-is-sqldev-093866.html>



NEW QUESTION 70

While creating the snapshots using the API, which Action should I be using?

- A. MakeSnapShot
- B. FreshSnapshot
- C. DeploySnapshot
- D. CreateSnapshot

Answer: D

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/CommandLineReference/ApiReference-cmd-CreateSnapshot.html>

NEW QUESTION 73

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

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

Answer: A

NEW QUESTION 78

What is the minimum charge for the data transferred between Amazon RDS and Amazon EC2 Instances in the same Availability Zone?

- A. USD 0.10 per GB
- B. No charge
- C. It is free.
- D. USD 0.02 per GB
- E. USD 0.01 per GB

Answer: B

Explanation:

For data transferred between an Amazon EC2 instance and Amazon RDS DB Instance in different Availability Zones of the same Region, there is no Data Transfer charge for traffic in or out of the Amazon RDS DB Instance. References:

NEW QUESTION 82

Which Amazon Storage behaves like raw, unformatted, external block devices that you can attach to your instances?

- A. None of these.
- B. Amazon Instance Storage
- C. Amazon EBS
- D. All of these

Answer: C

NEW QUESTION 87

What is the command line instruction for running the remote desktop client in Windows?

- A. desk.cpl
- B. mstsc

Answer: B

NEW QUESTION 89

In regards to IAM you can edit user properties later, but you cannot use the console to change the ____.

- A. user name
- B. password
- C. default group

Answer: A

NEW QUESTION 93

True or False: If you add a tag that has the same key as an existing tag on a DB Instance, the new value overwrites the old value.

- A. FALSE
- B. TRUE

Answer: B

Explanation:

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

NEW QUESTION 98

If I have multiple Read Replicas for my master DB Instance and I promote one of them, what happens to the rest of the Read Replicas?

- A. The remaining Read Replicas will still replicate from the older master DB Instance
- B. The remaining Read Replicas will be deleted
- C. The remaining Read Replicas will be combined to one read replica

Answer: A

Explanation:

If a source DB instance has several Read Replicas, promoting one of the Read Replicas to a DB instance has no effect on the other replicas.

NEW QUESTION 101

Because of the extensibility limitations of striped storage attached to Windows Server, Amazon RDS does not currently support increasing storage on a ____ DB Instance.

- A. SQL Server
- B. MySQL
- C. Oracle

Answer: A

NEW QUESTION 102

Through which of the following interfaces is AWS Identity and Access Management available?

- A) AWS Management Console
- B) Command line interface (CLI)
- C) IAM Query API
- D) Existing libraries

- A. Only through Command line interface (CLI)
- B. A, B and C
- C. A and C
- D. All of the above

Answer: D

Explanation:

Accessing IAM:

1 - AWS Management Console 2 - AWS Command Line Tools

3 - AWS SDKs (i.e. Existing libraries) 4 - IAM HTTPS API

<http://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html#intro-accessing>

NEW QUESTION 103

Select the incorrect statement

- A. In Amazon EC2, the private IP addresses only returned to Amazon EC2 when the instance is stopped or terminated
- B. In Amazon VPC, an instance retains its private IP addresses when the instance is stopped.
- C. In Amazon VPC, an instance does NOT retain its private IP addresses when the instance is stopped.
- D. In Amazon EC2, the private IP address is associated exclusively with the instance for its lifetime

Answer: C

Explanation:

A private IP address remains associated with the network interface when the instance is stopped and restarted, and is released when the instance is terminated.

NEW QUESTION 106

Can I delete a snapshot of the root device of an EBS volume used by a registered AMI?

- A. Only via API
- B. Only via Console
- C. Yes
- D. No

Answer: D

Explanation:

Note that you can't delete a snapshot of the root device of an EBS volume used by a registered AMI. You must first deregister the AMI before you can delete the snapshot.

Source: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-deleting-snapshot.html>

NEW QUESTION 107

The ____ service is targeted at organizations with multiple users or systems that use AWS products such as Amazon EC2, Amazon SimpleDB, and the AWS Management Console.

- A. Amazon RDS
- B. AWS Integrity Management
- C. AWS Identity and Access Management
- D. Amazon EMR

Answer: C

Explanation:

https://aws.amazon.com/documentation/iam/?nc1=h_ls

NEW QUESTION 110

True or False: Without IAM, you cannot control the tasks a particular user or system can do and what AWS resources they might use.

- A. FALSE
- B. TRUE

Answer: B

Explanation:

<http://docs.aws.amazon.com/IAM/latest/UserGuide/getting-setup.html>

NEW QUESTION 112

When automatic failover occurs, Amazon RDS will emit a DB Instance event to inform you that automatic failover occurred. You can use the ____ to return information about events related to your DB Instance

- A. FetchFailure
- B. DescribeFailure
- C. DescribeEvents
- D. FetchEvents

Answer: C

Explanation:

Q: Will I be alerted when automatic failover occurs?

Yes, Amazon RDS will emit a DB Instance event to inform you that automatic failover occurred. You can use the DescribeEvents to return information about events related to your DB Instance, or click the "DB Events" section of the AWS Management Console

<https://aws.amazon.com/rds/faqs/>

NEW QUESTION 113

How many Elastic IP by default in Amazon Account?

- A. 1 Elastic IP
- B. 3 Elastic IP
- C. 5 Elastic IP
- D. 0 Elastic IP

Answer: C

Explanation:

"By default, all AWS accounts are limited to 5 Elastic IP addresses, because public (IPv4) Internet addresses are a scarce public resource."

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html>

NEW QUESTION 114

The one-time payment for Reserved Instances is ____ refundable if the reservation is cancelled.

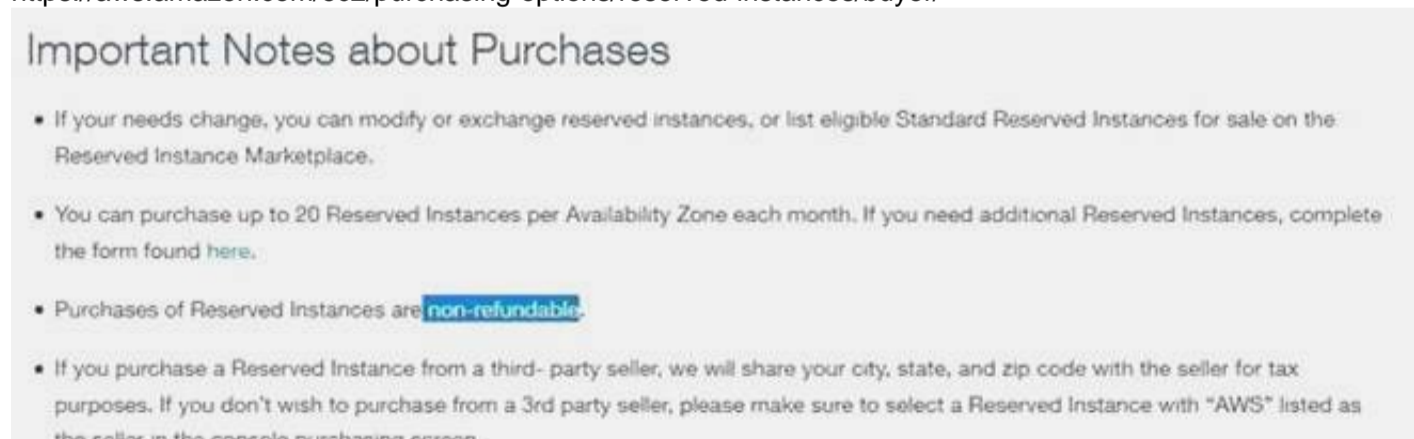
- A. always
- B. in some circumstances
- C. never

Answer: C

Explanation:

the one-time fee is non-refundable.

<https://aws.amazon.com/ec2/purchasing-options/reserved-instances/buyer/>



NEW QUESTION 119

If an Amazon EBS volume is the root device of an instance, can I detach it without stopping the instance?

- A. Yes but only if Windows instance
- B. No
- C. Yes
- D. Yes but only if a Linux instance

Answer: B

Explanation:

"If an EBS volume is the root device of an instance, you must stop the instance before you can detach the volume."
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-detaching-volume.html>

NEW QUESTION 122

A _____ is an individual, system, or application that interacts with AWS programmatically.

- A. user
- B. AWS Account
- C. Group
- D. Role

Answer: A

Explanation:

Q: What is a user?

A user is a unique identity recognized by AWS services and applications. Similar to a login user in an operating system like Windows or UNIX, a user has a unique name and can identify itself using familiar security credentials such as a password or access key. A user can be an individual, system, or application requiring access to AWS services. IAM supports users (referred to as "IAM users") managed in AWS's identity management system, and it also enables you to grant access to AWS resources for users managed outside of AWS in your corporate directory (referred to as "federated users").

NEW QUESTION 123

Can I initiate a "forced failover" for my MySQL Multi-AZ DB Instance deployment?

- A. Only in certain regions
- B. Only in VPC
- C. Yes
- D. No

Answer: C

Explanation:

If your DB instance is a Multi-AZ deployment, you can force a failover from one availability zone to another when you select the Reboot option. When you force a failover of your DB instance, Amazon RDS automatically switches to a standby replica in another Availability Zone and updates the DNS record for the DB instance to point to the standby DB instance. As a result, you will need to clean up and re-establish any existing connections to your DB instance. Reboot with failover is beneficial when you want to simulate a failure of a DB instance for testing, or restore operations to the original AZ after a failover occurs.

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

NEW QUESTION 125

A group can contain many users. Can a user belong to multiple groups?

- A. Yes always
- B. No
- C. Yes but only if they are using two factor authentication
- D. Yes but only in VPC

Answer: A

Explanation:

A group can contain many users, and a user can belong to multiple groups. http://docs.aws.amazon.com/IAM/latest/UserGuide/id_groups.html

NEW QUESTION 128

Please select the most correct answer regarding the persistence of the Amazon Instance Store

- A. The data on an instance store volume persists only during the life of the associated Amazon EC2 instance
- B. The data on an instance store volume is lost when the security group rule of the associated instance is changed.
- C. The data on an instance store volume persists even after associated Amazon EC2 instance is deleted

Answer: A

Explanation:

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



NEW QUESTION 133

My Read Replica appears "stuck" after a Multi-AZ failover and is unable to obtain or apply updates from the source DB Instance. What do I do?

- A. You will need to delete the Read Replica and create a new one to replace it.
- B. You will need to disassociate the DB Engine and CK associate it.
- C. The instance should be deployed to Single AZ and then moved to Multi- AZ once again
- D. You will need to delete the DB Instance and create a new one to replace i

Answer: A

Explanation:

Q: My Amazon RDS for MySQL Read Replica appears "stuck" after a Multi-AZ failover and is unable to obtain or apply updates from the source DB Instance. What do I do? ... To resolve the current issue, you will need to delete the Read Replica and create a new one to replace it. " <https://aws.amazon.com/rds/faqs/>

NEW QUESTION 137

Which DNS name can only be resolved within Amazon EC2?

- A. Internal DNS name
- B. External DNS name
- C. Global DNS name
- D. Private DNS name

Answer: D

Explanation:

To view DNS hostnames for an instance using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose Instances.
3. Select your instance from the list.
4. In the details pane, the Public DNS (IPv4) and Private DNS fields display the DNS hostnames, if applicable.

NEW QUESTION 140

If your DB instance runs out of storage space or file system resources, its status will change to _____ and your DB Instance will no longer be available.

- A. storage-overflow
- B. storage-full
- C. storage-exceed
- D. storage-overage

Answer: B

Explanation:

<https://aws.amazon.com/ko/premiumsupport/knowledge-center/rds-out-of-storage/>

Short Description

When an RDS DB instance reaches the **STORAGE_FULL** state, there is **not enough space available** for performing basic operations, eventually preventing you from restarting or making connections to the instance.

NEW QUESTION 145

Is it possible to access your EBS snapshots?

- A. Yes, through the Amazon S3 APIs.
- B. Yes, through the Amazon EC2 APIs.
- C. No, EBS snapshots cannot be accessed; they can only be used to create a new EBS volume.
- D. EBS doesn't provide snapshot

Answer: B

Explanation:

https://aws.amazon.com/ebs/faqs/?nc1=h_ls

Q: Will I be able to access my snapshots using the regular Amazon S3 API? No, snapshots are only available through the Amazon EC2 API.

NEW QUESTION 148

Because of the extensibility limitations of striped storage attached to Windows Server, Amazon RDS does not currently support increasing storage on a _____ DB Instance.

- A. SQL Server
- B. MySQL
- C. Oracle

Answer: A

NEW QUESTION 153

Which Amazon storage do you think is the best for my database-style applications that frequently encounter many random reads and writes across the dataset?

- A. None of these.
- B. Amazon Instance Storage
- C. Any of these
- D. Amazon EBS

Answer: D

Explanation:

"Amazon EBS is particularly helpful for database-style applications that frequently encounter many random reads and writes across the data set."
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html>

NEW QUESTION 155

In the context of MySQL, version numbers are organized as MySQL version = X.Y.Z. What does X denote here?

- A. release level
- B. minor version
- C. version number
- D. major version

Answer: D

Explanation:

MySQL on Amazon RDS Versions

For MySQL, version numbers are organized as version = X.Y.Z. In Amazon RDS terminology, **X.Y denotes the major version**, and **Z is the minor version number**. For Amazon RDS implementations, a version change is considered major if the major version number changes—for example, going from version 5.6 to 5.7. A version change is considered minor if only the minor version number changes—for example, going from version 5.6.22 to 5.6.23.

Amazon RDS currently supports MySQL major versions 5.5, 5.6, and 5.7. MySQL minor version support varies by AWS Region. Use the following table to see what MySQL minor versions are supported in each AWS Region.

NEW QUESTION 156

It is advised that you watch the Amazon CloudWatch " _____ " metric (available via the AWS Management Console or Amazon Cloud Watch APIs) carefully and recreate the Read Replica should it fall behind due to replication errors.

- A. Write Lag
- B. Read Replica
- C. Replica Lag
- D. Single Replica

Answer: C

Explanation:

The amount of time a Read Replica DB instance lags behind the source DB instance. Applies to MySQL, MariaDB, and PostgreSQL Read Replicas.
<http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/rds-metricscollected.html>

ReplicaLag	The amount of time a Read Replica DB instance lags behind the source DB instance. Applies to MySQL, MariaDB, and PostgreSQL Read Replicas.
	Units: Seconds

NEW QUESTION 158

By default, what are ENIs that are automatically created and attached to instances using the EC2 console set to do when the attached instance terminates?

- A. Remain as is
- B. Terminate
- C. Hibernate
- D. Pause

Answer: B

Explanation:

By default, elastic network interfaces that are automatically created and attached to instances using the console are set to terminate when the instance terminates. However, network interfaces created using the command line interface aren't set to terminate when the instance terminates.

NEW QUESTION 159

Are you able to integrate a multi-factor token service with the AWS Platform?

- A. Yes, you can integrate private multi-factor token devices to authenticate users to the AWS platform.
- B. No, you cannot integrate multi-factor token devices with the AWS platform.
- C. Yes, using the AWS multi-factor token devices to authenticate users on the AWS platfor

Answer: C

Explanation:

Private MFA does not apply here.

Q. What is AWS MFA?

AWS multi-factor authentication (AWS MFA) provides an extra level of security that you can apply to your AWS environment. You can enable AWS MFA for your AWS account and for individual AWS Identity and Access Management (IAM) users you create under your account.

NEW QUESTION 161

Fill in the blanks: _____ is a durable, block-level storage volume that you can attach to a single, running Amazon EC2 instance.

- A. Amazon S3
- B. Amazon EBS
- C. None of these
- D. All of these

Answer: B

NEW QUESTION 166

Location of Instances is _____

- A. Regional
- B. based on Availability Zone
- C. Global

Answer: B

Explanation:

Regions and Availability Zones

Amazon EC2 is hosted in multiple locations world-wide. These locations are composed of regions and Availability Zones. Each region is a separate geographic area. Each region has multiple, isolated locations known as Availability Zones. Amazon EC2 provides you the ability to place resources, such as instances, and data in multiple locations. Resources aren't replicated across regions unless you do so specifically. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availabilityzones.html#concepts-regions-availability-zones>

```
$ aws ec2 describe-availability-zones --region us-east-1
{
  "AvailabilityZones": [
    {
      "State": "available",
      "RegionName": "us-east-1",
      "Messages": [],
      "ZoneName": "us-east-1b"
    },
    {
      "State": "available",
      "RegionName": "us-east-1",
      "Messages": [],
      "ZoneName": "us-east-1c"
    },
    {
      "State": "available",
      "RegionName": "us-east-1",
      "Messages": [],
      "ZoneName": "us-east-1d"
    }
  ]
}
```

NEW QUESTION 171

What does Amazon SES stand for?

- A. Simple Elastic Server
- B. Simple Email Service
- C. Software Email Solution
- D. Software Enabled Server

Answer: B

Explanation:

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

Amazon **Simple Email Service** (Amazon SES) is a cost-effective email service built on the reliable and scalable infrastructure that Amazon.com developed to serve its own customer base. With Amazon SES, you can send and receive email with no required minimum commitments – you pay as you go, and you only pay for what you use.

NEW QUESTION 175

Can I attach more than one policy to a particular entity?

- A. Yes always
- B. Only if within GovCloud
- C. No
- D. Only if within VPC

Answer: A

NEW QUESTION 180

Can I detach the primary (eth0) network interface when the instance is running or stopped?

- A. Yes, You can.
- B. N
- C. You cannot
- D. Depends on the state of the interface at the time

Answer: B

Explanation:

Each instance in a VPC has a default elastic network interface (the primary network interface, eth0) that is assigned a private IP address from the IP address range of your VPC. You cannot detach a primary network interface from an instance.

NEW QUESTION 182

What does the "Server Side Encryption" option on Amazon S3 provide?

- A. It provides an encrypted virtual disk in the Cloud.
- B. It doesn't exist for Amazon S3, but only for Amazon EC2.
- C. It encrypts the files that you send to Amazon S3, on the server side.
- D. It allows to upload files using an SSL endpoint, for a secure transfe

Answer: C

Explanation:

Server-side encryption is about protecting data at rest. Server-side encryption with Amazon S3- managed encryption keys (SSE-S3) employs strong multi-factor encryption.

Amazon S3 encrypts each object with a unique key. As an additional safeguard, it encrypts the key itself with a master key that it regularly rotates. Amazon S3 server-side encryption uses one of the strongest block ciphers available, 256-bit Advanced Encryption Standard (AES-256), to encrypt your data.

References:

NEW QUESTION 186

What does Amazon EBS stand for?

- A. Elastic Block Storage
- B. Elastic Business Server
- C. Elastic Blade Server
- D. Elastic Block Store

Answer: D

Explanation:

<https://aws.amazon.com/ebs/> Amazon Elastic Block Store (EBS)

Amazon Elastic Block Store (Amazon EBS) provides persistent block level storage volumes for use with Amazon EC2 instances in the AWS Cloud. Each Amazon EBS volume is automatically replicated within its Availability Zone to protect you from component failure, offering high availability and durability. Amazon EBS volumes offer the consistent and low- latency performance needed to run your workloads. With Amazon EBS, you can scale your usage up or down within minutes all while paying a low price for only what you provision.

Amazon Elastic Block **Store** (Amazon EBS) provides persistent block storage volumes for use with Amazon EC2 instances in the AWS Cloud. Each Amazon EBS volume is automatically replicated within its Availability Zone to protect you from component failure, offering high availability and durability. Amazon EBS volumes offer the consistent and low-latency performance needed to run your workloads. With Amazon EBS, you can scale your usage up or down within minutes – all while paying a low price for only what you provision.

NEW QUESTION 187

Do the system resources on the Micro instance meet the recommended configuration for Oracle?

- A. Yes completely
- B. Yes but only for certain situations
- C. Not in any circumstance

Answer: C

Explanation:

We recommend that you use db.t1.micro instances with Oracle to test setup and connectivity only; the system resources for a db.t1.micro instance do not meet the recommended configuration for Oracle. No Oracle options are supported on a db.t1.micro instance.

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.DBInstanceClass.html#Concepts.DBInstanceClasses.Previous>

NEW QUESTION 188

Will I be charged if the DB instance is idle?

- A. No
- B. Yes
- C. Only is running in GovCloud
- D. Only if running in VPC

Answer: B

NEW QUESTION 191

Which choice is a storage option supported by Amazon EC2?

- A. Amazon SNS store
- B. Amazon Instance Store
- C. Amazon AppStream store
- D. None of these

Answer: B

Explanation:

Amazon EC2 supports the following storage options: Amazon Elastic Block Store (Amazon EBS) Amazon EC2 Instance Store
Amazon Simple Storage Service (Amazon S3) <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Storage.html>

NEW QUESTION 195

Without _____, you must either create multiple AWS accounts-each with its own billing and subscriptions to AWS products-or your employees must share the security credentials of a single AWS account.

- A. Amazon RDS
- B. Amazon Glacier
- C. Amazon EMR
- D. Amazon IAM

Answer: D

NEW QUESTION 197

Amazon RDS supports SOAP only through _____.

- A. HTTP or HTTPS
- B. TCP/IP
- C. HTTP
- D. HTTPS

Answer: D

Explanation:

Amazon RDS supports SOAP only through HTTPS
<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/using-soap-api.html>

WSDL and Schema Definitions

You can access the Amazon Relational Database Service using the SOAP web services messaging protocol. This interface is described by a Web Services Description Language (WSDL) document, which defines the operations and security model for the particular service. The WSDL references an XML Schema document, which strictly defines the data types that might appear in SOAP requests and responses. For more information on WSDL and SOAP, see [Web Services References](#).

Note

Amazon RDS supports SOAP only through HTTPS.

NEW QUESTION 200

The Amazon EC2 web service can be accessed using the _____ web services messaging protocol. This interface is described by a Web Services Description Language (WSDL) document.

- A. SOAP
- B. DCOM
- C. CORBA
- D. XML-RPC

Answer: A

Explanation:

<http://docs.aws.amazon.com/AWSECommerceService/latest/DG/WSDLLocation.html>

WSDL and Schema Definitions

You can access the Amazon Relational Database Service using the SOAP web services messaging protocol. This interface is described by a Web Services Description Language (WSDL) document, which defines the operations and security model for the particular service. The WSDL references an XML Schema document, which strictly defines the data types that might appear in SOAP requests and responses. For more information on WSDL and SOAP, see [Web Services References](#).

Note

Amazon RDS supports SOAP only through HTTPS.

NEW QUESTION 204

When you resize the Amazon RDS DB instance, Amazon RDS will perform the upgrade during the next maintenance window. If you want the upgrade to be performed now, rather than waiting for the maintenance window, specify the _____ option.

- A. ApplyNow
- B. ApplySoon
- C. ApplyThis
- D. ApplyImmediately

Answer: D

Explanation:

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.DBInstance.Modifying.html>

NEW QUESTION 209

The SQL Server ____ feature is an efficient means of copying data from a source database to your DB Instance. It writes the data that you specify to a data file, such as an ASCII file.

- A. bulk copy
- B. group copy
- C. dual copy
- D. mass copy

Answer: A

Explanation:

The SQL Server bulk copy feature is an efficient means of copying data from a source database to your DB Instance. Bulk copy writes the data that you specify to a data file, such as an ASCII

file. You can then run bulk copy again to write the contents of the file to the destination DB Instance.

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/SQLServer.Procedural.Importing.html>

NEW QUESTION 210

When using consolidated billing there are two account types. What are they?

- A. Paying account and Linked account
- B. Parent account and Child account
- C. Main account and Sub account.
- D. Main account and Secondary account

Answer: A

Explanation:

You sign up for Consolidated Billing in the AWS Billing and Cost Management console, and designate your account as a payer account. Now your account can pay the charges of the other accounts, which are called linked accounts. The payer account and the accounts linked to it are called a Consolidated Billing account family. Source: <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/consolidated-billing.html>

NEW QUESTION 213

If you're unable to connect via SSH to your EC2 instance, which of the following should you check and possibly correct to restore connectivity?

- A. Adjust Security Group to permit egress traffic over TCP port 443 from your IP.
- B. Configure the IAM role to permit changes to security group settings.
- C. Modify the instance security group to allow ingress of ICMP packets from your IP.
- D. Adjust the instance's Security Group to permit ingress traffic over port 22 from your IP.
- E. Apply the most recently released Operating System security patches

Answer: D

Explanation:

In a VPC everything is allowed out by default. References:

NEW QUESTION 214

Which of the following are characteristics of Amazon VPC subnets? (Choose two.)

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

Answer: BD

Explanation:

Even though we know the right Answers it is sometimes good to know why the other Answers are wrong.

A. Is wrong because a subnet maps to a single AZ.

C. Is wrong because /28 is the smallest subnet, amazon takes first four and last addresses per subnet.

E. Is wrong because a private subnet needs a NAT appliance.

NEW QUESTION 215

You have decided to change the instance type for instances running in your application tier that is using Auto Scaling. In which area below would you change the instance type definition?

- A. Auto Scaling policy
- B. Auto Scaling group
- C. Auto Scaling tags
- D. Auto Scaling launch configuration

Answer: D

NEW QUESTION 218

A company wants to implement their website in a virtual private cloud (VPC). The web tier will use an Auto Scaling group across multiple Availability Zones (AZs). The database will use Multi-AZ RDS MySQL and should not be publicly accessible. What is the minimum number of subnets that need to be configured in the VPC?

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

Answer: D

Explanation:

Since multi-AZ RDS needs 2 private subnets to provide high availability and 2 public subnets are needed for ELB(web-tier) application.

Would use VPC with private (DB) and public (WEB) subnets: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.Scenarios.html Multi AZ requirement forces me to multiply subnets by two.

Reasons:

For DB: Your VPC must have at least one subnet in at least two of the Availability Zones in the region where you want to deploy your DB instance. A subnet is a segment of a VPC's IP address

range that you can specify and that lets you group instances based on your security and operational needs

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

For Web: After creating a VPC, you can add one or more subnets in each Availability Zone. Each subnet must reside entirely within one Availability Zone and cannot span zones http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

NEW QUESTION 222

For which of the following use cases are Simple Workflow Service (SWF) and Amazon EC2 an appropriate solution? (Choose two.)

- A. Using as an endpoint to collect thousands of data points per hour from a distributed fleet of sensors
- B. Managing a multi-step and multi-decision checkout process of an e-commerce website
- C. Orchestrating the execution of distributed and auditable business processes
- D. Using as an SNS (Simple Notification Service) endpoint to trigger execution of video transcoding jobs
- E. Using as a distributed session store for your web application

Answer: BC

Explanation:

<https://aws.amazon.com/swf/faqs/>

NEW QUESTION 227

When creation of an EBS snapshot is initiated, but not completed, the EBS volume:

- A. Can be used while the snapshot is in progress.
- B. Cannot be detached or attached to an EC2 instance until the snapshot completes
- C. Can be used in read-only mode while the snapshot is in progress.
- D. Cannot be used until the snapshot complete

Answer: A

Explanation:

Snapshots occur asynchronously; the point-in-time snapshot is created immediately, but the status of the snapshot is pending until the snapshot is complete (when all of the modified blocks have been transferred to Amazon S3), which can take several hours for large initial snapshots or subsequent snapshots where many blocks have changed. While it is completing, an in-progress snapshot is not affected by ongoing reads and writes to the volume.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html>

NEW QUESTION 230

What are characteristics of Amazon S3? (Choose two.)

- A. S3 allows you to store objects of virtually unlimited size.
- B. S3 offers Provisioned IOPS.
- C. S3 allows you to store unlimited amounts of data.
- D. S3 should be used to host a relational database.
- E. Objects are directly accessible via a UR

Answer: CE

NEW QUESTION 232

A customer needs to capture all client connection information from their load balancer every five minutes. The company wants to use this data for analyzing traffic patterns and troubleshooting their applications. Which of the following options meets the customer requirements?

- A. Enable AWS CloudTrail for the load balancer.
- B. Enable access logs on the load balancer.

- C. Install the Amazon CloudWatch Logs agent on the load balancer.
- D. Enable Amazon CloudWatch metrics on the load balance

Answer: B

Explanation:

Elastic Load Balancing access logs

The access logs for Elastic Load Balancing capture detailed information for all requests made to your load balancer and stores them as log files in the Amazon S3 bucket that you specify. Each log contains

details such as the time a request was received, the client's IP address, latencies, request path, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot your back-end applications. For more information, see [Monitor Your Load Balancer Using Elastic Load Balancing Access Logs](#).

NEW QUESTION 235

If you want to launch Amazon Elastic Compute Cloud (EC2) instances and assign each instance a predetermined private IP address you should:

- A. Launch the instance from a private Amazon Machine Image (AMI).
- B. Assign a group of sequential Elastic IP address to the instances.
- C. Launch the instances in the Amazon Virtual Private Cloud (VPC).
- D. Launch the instances in a Placement Group.
- E. Use standard EC2 instances since each instance gets a private Domain Name Service (DNS) already

Answer: C

Explanation:

Each instance in a VPC has a default network interface (eth0) that is assigned the primary private IP address.

NEW QUESTION 237

An Auto-Scaling group spans 3 AZs and currently has 4 running EC2 instances. When Auto Scaling needs to terminate an EC2 instance by default, AutoScaling will:

(Choose two.)

- A. Allow at least five minutes for Windows/Linux shutdown scripts to complete, before terminating the instance.
- B. Terminate the instance with the least active network connection
- C. If multiple instances meet this criterion, one will be randomly selected.
- D. Send an SNS notification, if configured to do so.
- E. Terminate an instance in the AZ which currently has 2 running EC2 instances.
- F. Randomly select one of the 3 AZs, and then terminate an instance in that AZ

Answer: CD

Explanation:

Auto Scaling determines whether there are instances in multiple Availability Zones. If so, it selects the Availability Zone with the most instances and at least one instance that is not protected from scale in. <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AutoScalingBehavior.InstanceTermination.html>

NEW QUESTION 238

In order to optimize performance for a compute cluster that requires low inter-node latency, which of the following feature should you use?

- A. Multiple Availability Zones
- B. AWS Direct Connect
- C. EC2 Dedicated Instances
- D. Placement Groups
- E. VPC private subnets

Answer: D

Explanation:

A placement group is a logical grouping of instances within a single Availability Zone. Using placement groups enables applications to participate in a low-latency, 10 Gigabits per second (Gbps) network. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html> References:

NEW QUESTION 241

Which of the following notification endpoints or clients are supported by Amazon Simple Notification Service? (Choose two.)

- A. Email
- B. CloudFront distribution
- C. File Transfer Protocol
- D. Short Message Service
- E. Simple Network Management Protocol

Answer: AD

NEW QUESTION 244

You are deploying an application to track GPS coordinates of delivery trucks in the United States. Coordinates are transmitted from each delivery truck once every three seconds. You need to design an architecture that will enable real-time processing of these coordinates from multiple consumers. Which service should you use to implement data ingestion?

- A. Amazon Kinesis
- B. AWS Data Pipeline

- C. Amazon AppStream
- D. Amazon Simple Queue Service

Answer: A

Explanation:

<https://aws.amazon.com/streaming-data/>

NEW QUESTION 248

A photo-sharing service stores pictures in Amazon Simple Storage Service (S3) and allows application sign-in using an OpenID Connect-compatible identity provider. Which AWS Security Token Service approach to temporary access should you use for the Amazon S3 operations?

- A. SAML-based Identity Federation
- B. Cross-Account Access
- C. AWS Identity and Access Management roles
- D. Web Identity Federation

Answer: D

Explanation:

Web identity federation - You can let users sign in using a well-known third party identity provider such as Login with Amazon, Facebook, Google, or any OpenID Connect (OIDC) 2.0 compatible provider. AWS STS web identity federation supports Login with Amazon, Facebook, Google, and any OpenID Connect (OIDC)-compatible identity provider.

NEW QUESTION 249

A company is deploying a two-tier, highly available web application to AWS. Which service provides durable storage for static content while utilizing lower Overall CPU resources for the web tier?

- A. Amazon EBS volume
- B. Amazon S3
- C. Amazon EC2 instance store
- D. Amazon RDS instance

Answer: B

NEW QUESTION 252

A customer is hosting their company website on a cluster of web servers that are behind a publicfacing load balancer. The customer also uses Amazon Route 53 to manage their public DNS. How should the customer configure the DNS zone apex record to point to the load balancer?

- A. Create an A record pointing to the IP address of the load balancer
- B. Create a CNAME record pointing to the load balancer DNS name.
- C. Create a CNAME record aliased to the load balancer DNS name.
- D. Create an A record aliased to the load balancer DNS name

Answer: D

NEW QUESTION 253

You manually launch a NAT AMI in a public subnet. The network is properly configured. Security groups and network access control lists are property configured. Instances in a private subnet can access the NAT. The NAT can access the Internet. However, private instances cannot access the Internet. What additional step is required to allow access from the private instances?

- A. Enable Source/Destination Check on the private Instances.
- B. Enable Source/Destination Check on the NAT instance.
- C. Disable Source/Destination Check on the private instances.
- D. Disable Source/Destination Check on the NAT instanc

Answer: D

Explanation:

Disabling Source/Destination Checks.

Each EC2 instance performs source/destination checks by default. This means that the instance must be the source or destination of any traffic it sends or receives. However, a NAT instance must be able to send and receive traffic when the source or destination is not itself. Therefore, you must disable source/destination checks on the NAT instance. You can disable the SrcDestCheck attribute for a NAT instance that's either running or stopped using the console or the command line. http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_NAT_Instance.html

NEW QUESTION 254

Which of the following approaches provides the lowest cost for Amazon Elastic Block Store snapshots while giving you the ability to fully restore data?

- A. Maintain two snapshots: the original snapshot and the latest incremental snapshot.
- B. Maintain a volume snapshot; subsequent snapshots will overwrite one another
- C. Maintain a single snapshot the latest snapshot is both Incremental and complete.
- D. Maintain the most current snapshot, archive the original and incremental to Amazon Glacie

Answer: C

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-deleting-snapshot.html>

NEW QUESTION 256

A US-based company is expanding their web presence into Europe. The company wants to extend their AWS infrastructure from Northern Virginia (us-east-1) into the Dublin (eu-west-1) region. Which of the following options would enable an equivalent experience for users on both continents?

- A. Use a public-facing load balancer per region to load-balance web traffic, and enable HTTP health checks.
- B. Use a public-facing load balancer per region to load-balance web traffic, and enable sticky sessions.
- C. Use Amazon Route 53, and apply a geolocation routing policy to distribute traffic across both regions.
- D. Use Amazon Route 53, and apply a weighted routing policy to distribute traffic across both regions.

Answer: C

Explanation:

Geolocation routing lets you choose the resources that serve your traffic based on the geographic location of your users, meaning the location from which DNS queries originate. For example, you might want all queries from Africa to be routed to a web server with an IP address of 192.0.2.111. Another possible use is for balancing load across endpoints in a predictable, easy-to-manage way, so that each user location is consistently routed to the same endpoint.

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html#routing-policyweighted>

NEW QUESTION 260

You work for a major news network in Europe. They have just released a new app which allows users to report on events as and when they happen using their mobile phone. Users are able to upload pictures from the app and then other users will be able to view these pics. Your organization expects this app to grow very quickly, essentially doubling its user base every month. The app uses S3 to store the media and you are expecting sudden and large increases in traffic to S3 when a major news event takes place (as people will be uploading content in huge numbers). You need to keep your storage costs to a minimum however and it does not matter if some objects are lost. Which storage media should you use to keep costs as low as possible?

- A. S3 - Infrequently Accessed Storage.
- B. S3 - Reduced Redundancy Storage (RRS).
- C. Glacier.
- D. S3 - Provisioned IOP

Answer: B

NEW QUESTION 261

Your company has decided to set up a new AWS account for test and dev purposes. They already use AWS for production, but would like a new account dedicated for test and dev so as to not accidentally break the production environment. You launch an exact replica of your production environment using a CloudFormation template that your company uses in production. However CloudFormation fails. You use the exact same CloudFormation template in production, so the failure is something to do with your new AWS account. The CloudFormation template is trying to launch 60 new EC2 instances in a single AZ. After some research you discover that the problem is;

- A. For all new AWS accounts there is a soft limit of 20 EC2 instances per region
- B. You should submit the limit increase form and retry the template after your limit has been increased.
- C. For all new AWS accounts there is a soft limit of 20 EC2 instances per availability zone
- D. You should submit the limit increase form and retry the template after your limit has been increased.
- E. You cannot launch more than 20 instances in your default VPC, instead reconfigure the CloudFormation template to provision the instances in a custom VPC.
- F. Your CloudFormation template is configured to use the parent account and not the new account
- G. Change the account number in the CloudFormation template and relaunch the template.

Answer: A

NEW QUESTION 266

You run an automobile reselling company that has a popular online store on AWS. The application sits behind an Auto Scaling group and requires new instances of the Auto Scaling group to identify their public and private IP addresses. How can you achieve this?

- A. By using Ipconfig for windows or Ifconfig for Linux.
- B. By using a cloud watch metric.
- C. Using a Curl or Get Command to get the latest meta-data from <http://169.254.169.254/latest/meta-data/>
- D. Using a Curl or Get Command to get the latest user-data from <http://169.254.169.254/latest/userdata/>

Answer: C

NEW QUESTION 269

By definition a public subnet within a VPC is one that

- A. In its routing table it has at least one route that uses an Internet Gateway (IGW).
- B. Has at least one route in its routing table that routes via a Network Address Translation (NAT) instance.
- C. Where the Network Access Control List (NACL) permitting outbound traffic to 0.0.0.0/0.
- D. Has had the public subnet check box ticked when setting up this subnet in the VPC console

Answer: A

NEW QUESTION 270

You work in the genomics industry and you process large amounts of genomic data using a nightly Elastic Map Reduce (EMR) job. This job processes a single 3 Tb file which is stored on S3. The EMR job runs on 3 on-demand core nodes and four on-demand task nodes. The EMR job is now taking longer than anticipated and you have been asked to advise how to reduce the completion time?

- A. Use four Spot Instances for the task nodes rather than four On-Demand instances.

- B. You should reduce the input split size in the MapReduce job configuration and then adjust the number of simultaneous mapper tasks so that more tasks can be processed at once.
- C. Store the file on Elastic File Service instead of S3 and then mount EFS as an independent volume for your core nodes.
- D. Configure an independent VPC in which to run the EMR jobs and then mount EFS as an independent volume for your core nodes.
- E. Enable termination protection for the job flo

Answer: B

NEW QUESTION 273

You are a systems administrator and you need to monitor the health of your production environment. You decide to do this using Cloud Watch, however you notice that you cannot see the health of every important metric in the default dash board. Which of the following metrics do you need to design a custom cloud watch metric for, when monitoring the health of your EC2 instances?

- A. CPU Usage
- B. Memory usage
- C. Disk read operations
- D. Network in
- E. Estimated charges

Answer: B

NEW QUESTION 274

Amazon S3 buckets in all other regions (other than US Standard) provide read-after-write consistency for PUTS of new objects.

- A. True
- B. False

Answer: A

NEW QUESTION 277

When you create new subnets within a custom VPC, by default they can communicate with each other, across availability zones.

- A. True
- B. False

Answer: A

NEW QUESTION 280

You are creating your own relational database on an EC2 instance and you need to maximize IOPS performance. What can you do to achieve this goal?

- A. Add a single additional volume to the EC2 instance with provisioned IOPS.
- B. Create the database on an S3 bucket.
- C. Add multiple additional volumes with provisioned IOPS and then create a RAID 0 stripe across those volumes.
- D. Attach the single volume to multiple EC2 instances so as to maximize performanc

Answer: C

NEW QUESTION 282

Placement Groups can be created across 2 or more Availability Zones.

- A. True
- B. False

Answer: B

Explanation:



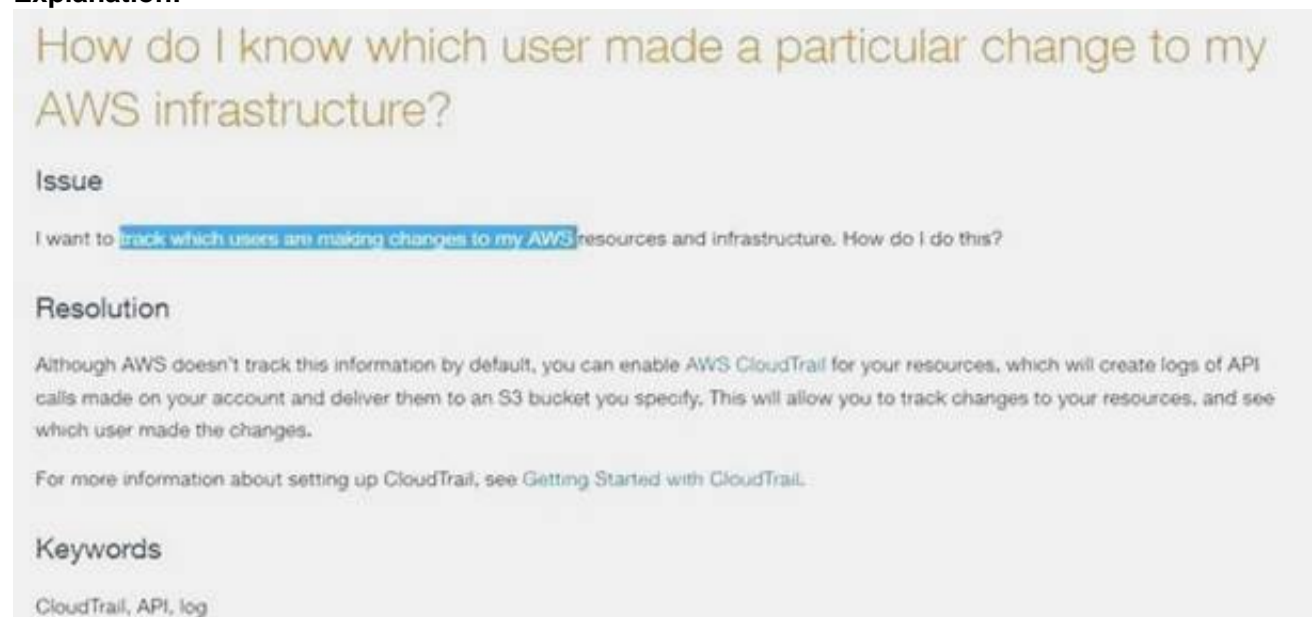
NEW QUESTION 286

You are appointed as your company's Chief Security Officer and you want to be able to track all changes made to your AWS environment, by all users and at all times, in all regions. What AWS service should you use to achieve this?

- A. CloudAudit
- B. CloudWatch
- C. CloudTrail
- D. CloudDetective

Answer: C

Explanation:



The screenshot shows an AWS support article. The title is "How do I know which user made a particular change to my AWS infrastructure?". Under the "Issue" section, it says "I want to track which users are making changes to my AWS resources and infrastructure. How do I do this?". Under the "Resolution" section, it explains that AWS CloudTrail can be enabled to track API calls and deliver logs to an S3 bucket. It also mentions that for more information, one should see "Getting Started with CloudTrail". Under the "Keywords" section, it lists "CloudTrail, API, log".

NEW QUESTION 291

You are designing a site for a new start up which generates cartoon images for people automatically. Customers will log on to the site, upload an image which is stored in S3. The application then passes a job to AWS SQS and a fleet of EC2 instances poll the queue to receive new processing jobs. These EC2 instances will then turn the picture in to a cartoon and will then need to store the processed job somewhere. Users will typically download the image once (immediately), and then never download the image again. What is the most commercially feasible method to store the processed images?

- A. Rather than use S3, store the images inside a BLOB on RDS with Multi-AZ configured for redundancy.
- B. Store the images on S3 RRS, and create a lifecycle policy to delete the image after 24 hours.
- C. Store the images on glacier instead of S3.
- D. Use elastic block storage volumes to store the image

Answer: B

NEW QUESTION 293

Which of the following is not a valid configuration type for AWS Storage gateway.

- A. Gateway-accessed volumes
- B. Gateway-cached volumes
- C. Gateway-stored volumes
- D. Gateway-Virtual Tape Library

Answer: A

NEW QUESTION 296

You work for a market analysis firm who are designing a new environment. They will ingest large amounts of market data via Kinesis and then analyze this data using Elastic Map Reduce. The data is then imported in to a high performance NoSQL Cassandra database which will run on EC2 and then be accessed by traders from around the world. The database volume itself will sit on 2 EBS volumes that will be grouped into a RAID 0 volume. They are expecting very high demand during peak times, with an IOPS performance level of approximately 15,000. Which EBS volume should you recommend?

- A. Magnetic
- B. General Purpose SSD
- C. Provisioned IOPS (PIOPS)
- D. Turbo IOPS (TIOPS)

Answer: C

Explanation:

Volume Type	Solid-State Drives (SSD)		Hard disk Drives (HDD)	
	General Purpose SSD (gp2)*	Provisioned IOPS SSD (io1)	Throughput Optimized HDD (st1)	Cold HDD (sc1)
Description	General purpose SSD volume that balances price and performance for a wide variety of transactional workloads	Highest-performance SSD volume designed for mission-critical applications	Low cost HDD volume designed for frequently accessed, throughput-intensive workloads	Lowest cost HDD volume designed for less frequently accessed workloads
Use Cases	<ul style="list-style-type: none"> Recommended for most workloads System boot volumes Virtual desktops Low-latency interactive apps Development and test environments 	<ul style="list-style-type: none"> Critical business applications that require sustained IOPS performance, or more than 10,000 IOPS or 160 MiB/s of throughput per volume Large database workloads, such as: <ul style="list-style-type: none"> MongoDB Cassandra Microsoft SQL Server MySQL PostgreSQL Oracle 	<ul style="list-style-type: none"> Streaming workloads requiring consistent, fast throughput at a low price Big data Data warehouses Log processing Cannot be a boot volume 	<ul style="list-style-type: none"> Throughput-oriented storage for large volumes of data that is infrequently accessed Scenarios where the lowest storage cost is important Cannot be a boot volume
API Name	gp2	io1	st1	sc1
Volume Size	1 GiB - 16 TiB	4 GiB - 16 TiB	500 GiB - 16 TiB	500 GiB - 16 TiB
Max. IOPS**/Volume	10,000	20,000	500	250

NEW QUESTION 301

You are a solutions architect working for a large digital media company. Your company is migrating their production estate to AWS and you are in the process of setting up access to the AWS console using Identity Access Management (IAM). You have created 5 users for your system administrators. What further steps do you need to take to enable your system administrators to get access to the AWS console?

- Generate an Access Key ID & Secret Access Key, and give these to your system administrators.
- Enable multi-factor authentication on their accounts and define a password policy.
- Generate a password for each user created and give these passwords to your system administrators.
- Give the system administrators the secret access key and access key id, and tell them to use these credentials to log in to the AWS console.

Answer: C

NEW QUESTION 302

You are a solutions architect working for a company that specializes in ingesting large data feeds (using Kinesis) and then analyzing these feeds using Elastic Map Reduce (EMR). The results are then stored on a custom MySQL database which is hosted on an EC2 instance which has 3 volumes, the root/boot volume, and then 2 additional volumes which are striped in to a RAID 1. Your company recently had an outage and lost some key data and have since decided that they will need to run nightly back ups. Your application is only used during office hours, so you can afford to have some down time in the middle of the night if required. You decide to take a snapshot of all three volumes every 24 hours. In what manner should you do this?

- Take a snapshot of each volume independently, while the EC2 instance is running.
- Stop the EC2 instance and take a snapshot of each EC2 instance independentl
- Once the snapshots are complete, start the EC2 instance and ensure that all relevant volumes are remounted.
- Add two additional volumes to the existing RAID 0 volume and mirror these volumes creating a RAID 10. Take a snap of only the two new volumes.
- Create a read replica of the existing EC2 instance and then take your snapshots from the read replica and not the live EC2 instance.

Answer: B

NEW QUESTION 307

Multi-AZ deployment is supported for Microsoft SQL Server DB Instances.

- True
- False

Answer: A

NEW QUESTION 312

While creating an EC2 snapshot using the API, which Action should I be using?

- MakeSnapshot
- FreshSnapshot

- C. DeploySnapshot
- D. CreateSnapshot

Answer: D

NEW QUESTION 315

SQL Server stores logins and passwords in the master database.

- A. True
- B. False

Answer: A

Explanation:

There are two authentications

Windows authentication

The credentials for which are not stored in SQL Server database and managed by windows/AD. There would be entry for windows authenticated logins in master database with respective SID but password would be with Active directory.

SQL Server authentication.

For 2nd we have password stored in hash format you can see it from sys.sql_logins. The information about SQL server logins are stored in master database and each login has SID respective to it. Only SA login has same SID no matter what server it is. That is why when you move database by backup restore mechanism users are moved not logins and you finally have to create logins(if already not there) and map it to users. This is generally called as troubleshooting orphaned users

NEW QUESTION 317

While performing volume status checks using volume status checks, if the status is insufficient-data, what does it mean?

- A. checks may still be in progress on the volume
- B. check has passed
- C. check has failed
- D. there is no such status

Answer: A

Explanation:

Volume status checks are automated tests that run every 5 minutes and return a pass or fail status. If all checks pass, the status of the volume is ok. If a check fails, the status of the volume is impaired. If the status is insufficient-data, the checks may still be in progress on the volume.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring-volume-status.html>

NEW QUESTION 322

What is the maximum write throughput I can provision per table for a single DynamoDB table?

- A. 5,000 us east, 1,000 all other regions
- B. 100,000 us east, 10, 000 all other regions
- C. Designed to scale without limits, but if you go beyond 40,000 us east/10,000 all other regions you have to contact AWS first.
- D. There is no limit

Answer: C

NEW QUESTION 323

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

- A. Raid 5
- B. Raid 6
- C. Raid 1
- D. Raid 2

Answer: C

NEW QUESTION 325

In the basic monitoring package for EC2, Amazon CloudWatch provides the following metrics:

- A. web server visible metrics such as number failed transaction requests
- B. operating system visible metrics such as memory utilization
- C. database visible metrics such as number of connections
- D. hypervisor visible metrics such as CPU utilization

Answer: D

NEW QUESTION 327

Which of the following will occur when an EC2 instance in a VPC with an associated Elastic IP is stopped and started? (Choose 2 answers)

- A. The Elastic IP will be dissociated from the instance
- B. All data on instance-store devices will be lost
- C. All data on EBS (Elastic Block Store) devices will be lost
- D. The ENI (Elastic Network Interface) is detached
- E. The underlying host for the instance is changed

Answer: BE

NEW QUESTION 332

Which route must be added to your routing table in order to allow connections to the Internet from your subnet?

- A. Destination: 0.0.0.0/0 --> Target: your Internet gateway
- B. Destination: 192.168.1.257/0 --> Target: your Internet gateway
- C. Destination: 0.0.0.0/33 --> Target: your virtual private gateway
- D. Destination: 0.0.0.0/0 --> Target: 0.0.0.0/24
- E. Destination: 10.0.0.0/32 --> Target: your virtual private gateway

Answer: A

NEW QUESTION 333

You are developing a highly available web application using stateless web servers. Which services are suitable for storing session state data? (Choose three.)

- A. Amazon CloudWatch
- B. Amazon Relational Database Service (RDS)
- C. Elastic Load Balancing
- D. Amazon ElastiCache
- E. AWS Storage Gateway
- F. Amazon DynamoDB

Answer: BDF

NEW QUESTION 335

A VPC public subnet is one that:

- A. Has at least one route in its associated routing table that uses an Internet Gateway (IGW).
- B. Includes a route in its associated routing table via a Network Address Translation (NAT) instance.
- C. Has a Network Access Control List (NACL) permitting outbound traffic to 0.0.0.0/0.
- D. Has the Public Subnet option selected in its configuratio

Answer: A

Explanation:

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

If a subnet's traffic is routed to an Internet gateway, the subnet is known as a public subnet.

NEW QUESTION 339

A customer has a requirement to extend their on-premises data center to AWS. The customer requires a 50-Mbps dedicated and private connection to their VPC. Which AWS product or feature satisfies this requirement?

- A. Amazon VPC Peering
- B. Elastic IP Addresses
- C. AWS Direct Connect
- D. Amazon VPC virtual private gateway

Answer: C

Explanation:

AWS Direct connect is the solution officially provided by AWS when the customer wants to have a dedicated and private connection to their AWS cloud.

The correct answer is: **AWS Direct Connect**

NEW QUESTION 343

You have a web portal composed of two services. Each service must scale independently. Both services should be served under the same domain. Which configuration allows this?

- A. Use two AWS Application Load Balancers: one for each service
- B. Assign the same CNAME to both.
- C. Use one AWS Classic Load Balance
- D. Create a redirect in the web server based on user's source IPs.

- E. Use two AWS Classic Load Balancers: one for each service
- F. Assign the same CNAME to both.
- G. Use one AWS Application Load Balance
- H. Specify listener rules to route requests to each service

Answer: C

NEW QUESTION 344

Which Auto Scaling features allow you to scale ahead of expected increases in load? (Select TWO.)

- A. Cooldown period
- B. Lifecycle hooks
- C. Desired capacity
- D. Scheduled scaling
- E. Health check grace period
- F. Metric-based scaling

Answer: DF

NEW QUESTION 347

You are designing a scalable web application with stateless web servers. Which service or feature is well suited to store user session information?

- A. Amazon EBS
- B. Amazon DynamoDB
- C. Amazon EC2 instance store
- D. Amazon SQS

Answer: C

NEW QUESTION 352

Your Amazon EC2 instances must access the AWS API, so you created a NAT gateway in an existing subnet. When you try to access the AWS API, you are unsuccessful.

What could be preventing access?

- A. The NAT gateway subnet does not have a route to an Internet gateway.
- B. The instances need an IAM granting access to the NAT gateway.
- C. The NAT gateway does not have a route to the virtual private gateway.
- D. The instances are not in the same subnet as the NAT gateway.

Answer: A

NEW QUESTION 357

Your existing web application requires a persistent key-value store database that must service 50,000 reads/second. Your company is looking at 10% growth in traffic and data volume month over month for the next several years.

Which service meets these requirements?

- A. Amazon Redshift
- B. Amazon DynamoDB
- C. Amazon SQS
- D. Amazon RDS

Answer: D

NEW QUESTION 358

Your company's IT policies mandate that all critical data must be duplicated in two physical locations at least 100 miles apart.

Which storage option meets this requirement?

- A. Two Amazon S3 buckets in different regions
- B. One Amazon S3 bucket
- C. One Amazon Glacier archive
- D. Two Amazon S3 buckets in the same region

Answer: A

NEW QUESTION 362

A colleague asked for your advice about how to easily deploy, monitor, and scale a three-tier LAMP (Linux, Apache, MySQL, PHP) application on AWS. Your colleague has time and staffing constraints and wants to deploy and manage the application with minimal effort.

Which AWS service would you suggest?

- A. Elastic Beanstalk
- B. Data Pipeline
- C. CloudFormation
- D. CodeDeploy

Answer: A

NEW QUESTION 365

Which services can invoke AWS Lambda functions? (Select TWO.)

- A. Amazon SNS
- B. Amazon Redshift
- C. Amazon Route53
- D. Amazon DynamoDB
- E. Elastic Load Balancing

Answer: AD

NEW QUESTION 369

A customer's security team requires the logging of all network access attempts to Amazon EC2 instances in their production VPC on AWS. Which configuration will meet the security team's requirement?

- A. Enable CloudTrail for the production VPC.
- B. Enable VPC Flow Logs for the production VPC.
- C. Enable both CloudTrail and VPC Flow Logs for the production VPC.
- D. Enable both CloudTrail and VPC Flow Logs for the AWS account.

Answer: B

NEW QUESTION 371

Your company runs an application that generates several thousand 1-GB reports a month. Approximately 10% of these reports will be accessed once during the first 30 days and must be available on demand. After 30 days, reports are no longer accessed as a part of normal business processes but must be retained for compliance reasons.

Which architecture would meet these requirements with the lowest cost?

- A. Upload the reports to Amazon S3 Standard storage class
- B. Set a lifecycle configuration on the bucket to transition the reports to Amazon Glacier after 30 days.
- C. Upload the reports to Amazon S3 Standard – Infrequent Access storage class
- D. Set a lifecycle configuration on the bucket to transition the reports to Amazon Glacier after 30 days.
- E. Upload the reports to Amazon Glacier
- F. When reports are requested, copy them to Amazon S3 Standard storage class for access
- G. Delete the copied reports after they have been viewed.
- H. Upload the reports to Amazon S3 Standard – Infrequent Access storage class
- I. When reports are requested, copy them to Amazon S3 Standard storage class for access
- J. Delete the copied reports after they have been viewed.

Answer: B

NEW QUESTION 375

A stray Amazon EC2 r3.xlarge instance is running in your AWS account. Before terminating it, you want to find the owner to confirm that it is not needed. Where can you find the identity that launched this instance?

- A. VPC flow logs
- B. ELB access logs
- C. CloudTrail logs
- D. Operating system logs

Answer: C

NEW QUESTION 376

You are running a mobile media application and are considering API Gateway for the client entry point. What benefits would this provide? (Select TWO.)

- A. Caching API responses
- B. IP blacklisting
- C. Intrusion prevention
- D. Load balancing
- E. Throttling traffic

Answer: AE

NEW QUESTION 381

Amazon EBS provides the ability to create backups of any Amazon EC2 volume into what is known as _____.

- A. snapshots
- B. images
- C. instance backups
- D. mirrors

Answer: A

Explanation:

Amazon allows you to make backups of the data stored in your EBS volumes through snapshots that can later be used to create a new EBS volume.

NEW QUESTION 383

To specify a resource in a policy statement, in Amazon EC2, can you use its Amazon Resource Name (ARN)?

- A. Yes, you can.
- B. No, you can't because EC2 is not related to ARN.
- C. No, you can't because you can't specify a particular Amazon EC2 resource in an IAM policy.
- D. Yes, you can but only for the resources that are not affected by the action.

Answer: A

Explanation:

Some Amazon EC2 API actions allow you to include specific resources in your policy that can be created or modified by the action. To specify a resource in the statement, you need to use its Amazon Resource Name (ARN).

NEW QUESTION 387

One of the criteria for a new deployment is that the customer wants to use AWS Storage Gateway. However, you are not sure whether you should use gateway-cached volumes or gateway-stored volumes or even what the differences are. Which statement below best describes those differences?

- A. Gateway-cached lets you store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally.
- B. Gateway-stored enables you to configure your on-premises gateway to store all your data locally and then asynchronously back up point-in-time snapshots of this data to Amazon S3.
- C. Gateway-cached is free whilst gateway-stored is not.
- D. Gateway-cached is up to 10 times faster than gateway-stored.
- E. Gateway-stored lets you store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally.
- F. Gateway-cached enables you to configure your on-premises gateway to store all your data locally and then asynchronously back up point-in-time snapshots of this data to Amazon S3.

Answer: A

Explanation:

Volume gateways provide cloud-backed storage volumes that you can mount as Internet Small Computer System Interface (iSCSI) devices from your on-premises application servers. The gateway supports the following volume configurations:

Gateway-cached volumes. You store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally.

Gateway-cached volumes offer a substantial cost savings on primary storage and minimize the need to scale your storage on-premises. You also retain low-latency access to your frequently accessed data.

Gateway-stored volumes. If you need low-latency access to your entire data set, you can configure your on-premises gateway to store all your data locally and then asynchronously back up point-in-time snapshots of this data to Amazon S3. This configuration provides durable and inexpensive offsite backups that you can recover to your local data center or Amazon EC2. For example, if you need replacement capacity for disaster recovery, you can recover the backups to Amazon EC2. References:

NEW QUESTION 390

What is a placement group in Amazon EC2?

- A. It is a group of EC2 instances within a single Availability Zone.
- B. It is the edge location of your web content.
- C. It is the AWS region where you run the EC2 instance of your web content.
- D. It is a group used to span multiple Availability Zones.

Answer: A

Explanation:

A placement group is a logical grouping of instances within a single Availability Zone. References:

NEW QUESTION 392

A client needs you to import some existing infrastructure from a dedicated hosting provider to AWS to try and save on the cost of running his current website. He also needs an automated process that manages backups, software patching, automatic failure detection, and recovery. You are aware that his existing set up currently uses an Oracle database. Which of the following AWS databases would be best for accomplishing this task?

- A. Amazon RDS
- B. Amazon Redshift
- C. Amazon SimpleDB
- D. Amazon ElastiCache

Answer: A

Explanation:

Amazon RDS gives you access to the capabilities of a familiar MySQL, Oracle, SQL Server, or PostgreSQL database engine. This means that the code, applications, and tools you already use today with your existing databases can be used with Amazon RDS. Amazon RDS automatically patches the database software and backs up your database, storing the backups for a user-defined retention period and enabling point-in-time recovery.

References:

NEW QUESTION 395

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.

NEW QUESTION 397

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.

NEW QUESTION 398

In Amazon AWS, which of the following statements is true of key pairs?

- A. Key pairs are used only for Amazon SDKs.
- B. Key pairs are used only for Amazon EC2 and Amazon CloudFront.
- C. Key pairs are used only for Elastic Load Balancing and AWS IAM.
- D. Key pairs are used for all Amazon service

Answer: B

Explanation:

Key pairs consist of a public and private key, where you use the private key to create a digital signature, and then AWS uses the corresponding public key to validate the signature. Key pairs are used only for Amazon EC2 and Amazon CloudFront.

NEW QUESTION 399

An organization has three separate AWS accounts, one each for development, testing, and production. The organization wants the testing team to have access to certain AWS resources in the production account. How can the organization achieve this?

- A. It is not possible to access resources of one account with another account.
- B. Create the IAM roles with cross account access.
- C. Create the IAM user in a test account, and allow it access to the production environment with the IAM policy.
- D. Create the IAM users with cross account access

Answer: B

Explanation:

An organization has multiple AWS accounts to isolate a development environment from a testing or production environment. At times the users from one account need to access resources in the other account, such as promoting an update from the development environment to the production environment. In this case the IAM role with cross account access will provide a solution. Cross account access lets one account share access to their resources with users in the other AWS accounts.

NEW QUESTION 402

You need to import several hundred megabytes of data from a local Oracle database to an Amazon RDS DB instance. What does AWS recommend you use to accomplish this?

- A. Oracle export/import utilities
- B. Oracle SQL Developer
- C. Oracle Data Pump
- D. DBMS_FILE_TRANSFER

Answer: C

Explanation:

How you import data into an Amazon RDS DB instance depends on the amount of data you have and the number and variety of database objects in your database.

For example, you can use Oracle SQL Developer to import a simple, 20 MB database; you want to use Oracle Data Pump to import complex databases or databases that are several hundred megabytes or several terabytes in size.

NEW QUESTION 407

A user has created an EBS volume with 1000 IOPS. What is the average IOPS that the user will get for most of the year as per EC2 SLA if the instance is attached to the EBS optimized instance?

- A. 950

- B. 990
- C. 1000
- D. 900

Answer: D

Explanation:

As per AWS SLA if the instance is attached to an EBS-Optimized instance, then the Provisioned IOPS volumes are designed to deliver within 10% of the provisioned IOPS performance 99.9% of the time in a given year. Thus, if the user has created a volume of 1000 IOPS, the user will get a minimum 900 IOPS 99.9% time of the year.

NEW QUESTION 408

You are in the process of creating a Route 53 DNS failover to direct traffic to two EC2 zones. Obviously, if one fails, you would like to direct Route 53 traffic to the other region. Each region has an ELB with some instances being distributed. What is the best way for you to configure the Route 53 health check?

- A. Route 53 doesn't support ELB with an internal health check
- B. You need to create your own Route 53 health check of the ELB
- C. Route 53 natively supports ELB with an internal health check
- D. Turn "Evaluate target health" off and "Associate with Health Check" on and Route 53 will use the ELB's internal health check.
- E. Route 53 doesn't support ELB with an internal health check
- F. You need to associate your resource record set for the ELB with your own health check
- G. Route 53 natively supports ELB with an internal health check
- H. Turn "Evaluate target health" on and "Associate with Health Check" off and Route 53 will use the ELB's internal health check.

Answer: D

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. When you enable this feature, Route 53 uses health checks--regularly making Internet requests to your application's endpoints from multiple locations around the world--to determine whether each endpoint of your application is up or down. To enable DNS Failover for an ELB endpoint, create an Alias record pointing to the ELB and set the "Evaluate Target Health" parameter to true. Route 53 creates and manages the health checks for your ELB automatically. You do not need to create your own Route 53 health check of the ELB. You also do not need to associate your resource record set for the ELB with your own health check, because Route 53 automatically associates it with the health checks that Route 53 manages on your behalf. The ELB health check will also inherit the health of your backend instances behind that ELB.

NEW QUESTION 412

A user wants to use an EBS-backed Amazon EC2 instance for a temporary job. Based on the input data, the job is most likely to finish within a week. Which of the following steps should be followed to terminate the instance automatically once the job is finished?

- A. Configure the EC2 instance with a stop instance to terminate it.
- B. Configure the EC2 instance with ELB to terminate the instance when it remains idle.
- C. Configure the CloudWatch alarm on the instance that should perform the termination action once the instance is idle.
- D. Configure the Auto Scaling schedule activity that terminates the instance after 7 days

Answer: C

Explanation:

Auto Scaling can start and stop the instance at a pre-defined time. Here, the total running time is unknown. Thus, the user has to use the CloudWatch alarm, which monitors the CPU utilization.

The user can create an alarm that is triggered when the average CPU utilization percentage has been lower than 10 percent for 24 hours, signaling that it is idle and no longer in use. When the utilization is below the threshold limit, it will terminate the instance as a part of the instance action. References:

NEW QUESTION 417

While using the EC2 GET requests as URLs, the is the URL that serves as the entry point for the web service.

- A. token
- B. endpoint
- C. action
- D. None of these

Answer: B

Explanation:

The endpoint is the URL that serves as the entry point for the web service.

NEW QUESTION 418

You are building infrastructure for a data warehousing solution and an extra request has come through that there will be a lot of business reporting queries running all the time and you are not sure if your current DB instance will be able to handle it. What would be the best solution for this?

- A. DB Parameter Groups
- B. Read Replicas
- C. Multi-AZ DB Instance deployment
- D. Database Snapshots

Answer: B

Explanation:

Read Replicas make it easy to take advantage of MySQL's built-in replication functionality to elastically scale out beyond the capacity constraints of a single DB

Instance for read-heavy database workloads. There are a variety of scenarios where deploying one or more Read Replicas for a given source DB Instance may make sense. Common reasons for deploying a Read Replica include: Scaling beyond the compute or I/O capacity of a single DB Instance for read-heavy database workloads. This excess read traffic can be directed to one or more Read Replicas. Serving read traffic while the source DB Instance is unavailable. If your source DB Instance cannot take I/O requests (e.g. due to I/O suspension for backups or scheduled maintenance), you can direct read traffic to your Read Replica(s). For this use case, keep in mind that the data on the Read Replica may be "stale" since the source DB Instance is unavailable. Business reporting or data warehousing scenarios; you may want business reporting queries to run against a Read Replica, rather than your primary, production DB Instance.

NEW QUESTION 419

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.

NEW QUESTION 422

Which of the following AWS CLI commands is syntactically incorrect?

- 1. `$ aws ec2 describe-instances`
- 2. `$ aws ec2 start-instances --instance-ids i-1348636c`
- 3. `$ aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError -message "Script Failure"`
- 4. `$ aws sqs receive-message --queue-url https://queue.amazonaws.com/546419318123/Test`

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

Answer: A

Explanation:

The following CLI command is missing a hyphen before "-message". `aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError -message "Script Failure"`.

It has been added below in red `aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError ---message "Script Failure"`

NEW QUESTION 425

A customer enquires about whether all his data is secure on AWS and is especially concerned about Elastic Map Reduce (EMR) so you need to inform him of some of the security features in place for AWS. Which of the below statements would be an incorrect response to your customers enquiry?

- A. Amazon EMR customers can choose to send data to Amazon S3 using the HTTPS protocol for secure transmission.
- B. Amazon S3 provides authentication mechanisms to ensure that stored data is secured against unauthorized access.
- C. Every packet sent in the AWS network uses Internet Protocol Security (IPsec).
- D. Customers may encrypt the input data before they upload it to Amazon S3.

Answer: C

Explanation:

Amazon S3 provides authentication mechanisms to ensure that stored data is secured against unauthorized access. Unless the customer who is uploading the data specifies otherwise, only that customer can access the data. Amazon EMR customers can also choose to send data to Amazon S3 using the HTTPS protocol for secure transmission. In addition, Amazon EMR always uses HTTPS to send data between Amazon S3 and Amazon EC2. For added security, customers may encrypt the input data before they upload it to Amazon S3 (using any common data compression tool); they then need to add a decryption step to the beginning of their cluster when Amazon EMR fetches the data from Amazon S3.

NEW QUESTION 429

A major customer has asked you to set up his AWS infrastructure so that it will be easy to recover in the case of a disaster of some sort. Which of the following is important when thinking about being able to quickly launch resources in AWS to ensure business continuity in case of a disaster?

- A. Create and maintain AMIs of key servers where fast recovery is required.
- B. Regularly run your servers, test them, and apply any software updates and configuration changes.
- C. All items listed here are important when thinking about disaster recovery.
- D. Ensure that you have all supporting custom software packages available in AWS

Answer: C

Explanation:

In the event of a disaster to your AWS infrastructure you should be able to quickly launch resources in Amazon Web Services (AWS) to ensure business continuity.

The following are some key steps you should have in place for preparation:

- 1. Set up Amazon EC2 instances to replicate or mirror data.
- 2. Ensure that you have all supporting custom software packages available in AWS.
- 3. Create and maintain AMIs of key servers where fast recovery is required.
- 4. Regularly run these servers, test them, and apply any software updates and configuration changes.
- 5. Consider automating the provisioning of AWS resources. References:

NEW QUESTION 432

You need to measure the performance of your EBS volumes as they seem to be under performing. You have come up with a measurement of 1,024 KB I/O but your colleague tells you that EBS volume performance is measured in IOPS. How many IOPS is equal to 1,024 KB I/O?

- A. 16
- B. 256
- C. 8
- D. 4

Answer: D

Explanation:

Several factors can affect the performance of Amazon EBS volumes, such as instance configuration, I/O characteristics, workload demand, and storage configuration. IOPS are input/output operations per second. Amazon EBS measures each I/O operation per second (that is 256 KB or smaller) as one IOPS. I/O operations that are larger than 256 KB are counted in 256 KB capacity units.

For example, a 1,024 KB I/O operation would count as 4 IOPS. When you provision a 4,000 IOPS volume and attach it to an EBS-optimized instance that can provide the necessary bandwidth, you can transfer up to 4,000 chunks of data per second (provided that the I/O does not exceed the 128 MB/s per volume throughput limit of General Purpose (SSD) and Provisioned IOPS (SSD) volumes).

NEW QUESTION 437

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.

NEW QUESTION 442

An existing client comes to you and says that he has heard that launching instances into a VPC (virtual private cloud) is a better strategy than launching instances into a EC2-classic which he knows is what you currently do. You suspect that he is correct and he has asked you to do some research about this and get back to him. Which of the following statements is true in regards to what ability launching your instances into a VPC instead of EC2-Classical gives you?

- A. All of the things listed here.
- B. Change security group membership for your instances while they're running
- C. Assign static private IP addresses to your instances that persist across starts and stops
- D. Define network interfaces, and attach one or more network interfaces to your instances

Answer: A

Explanation:

By launching your instances into a VPC instead of EC2-Classical, you gain the ability to:

Assign static private IP addresses to your instances that persist across starts and stops Assign multiple IP addresses to your instances.

Define network interfaces, and attach one or more network interfaces to your instances Change security group membership for your instances while they're running Control the outbound traffic from your instances (egress filtering) in addition to controlling the inbound traffic to them (ingress filtering).

Add an additional layer of access control to your instances in the form of network access control lists (ACL).

Run your instances on single-tenant hardware

NEW QUESTION 444

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.

NEW QUESTION 449

You have set up an Elastic Load Balancer (ELB) with the usual default settings, which route each request independently to the application instance with the

smallest load. However, someone has asked you to bind a user's session to a specific application instance so as to ensure that all requests coming from the user during the session will be sent to the same application instance. AWS has a feature to do this. What is it called?

- A. Connection draining
- B. Proxy protocol
- C. Tagging
- D. Sticky session

Answer: D

Explanation:

An Elastic Load Balancer(ELB) by default, routes each request independently to the application instance with the smallest load. However, you can use the sticky session feature (also known as session affinity), which enables the load balancer to bind a user's session to a specific application instance. This ensures that all requests coming from the user during the session will be sent to the same application instance. The key to managing the sticky session is determining how long your load balancer should consistently route the user's request to the same application instance. If your application has its own session cookie, then you can set Elastic Load Balancing to create the session cookie to follow the duration specified by the application's session cookie. If your application does not have its own session cookie, then you can set Elastic Load Balancing to create a session cookie by specifying your own stickiness duration. You can associate stickiness duration for only HTTP/HTTPS load balancer listeners. An application instance must always receive and send two cookies: A cookie that defines the stickiness duration and a special Elastic Load Balancing cookie named AWSELB, that has the mapping to the application instance.

NEW QUESTION 453

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SAA-C01 Practice Exam Features:

- * SAA-C01 Questions and Answers Updated Frequently
- * SAA-C01 Practice Questions Verified by Expert Senior Certified Staff
- * SAA-C01 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * SAA-C01 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SAA-C01 Practice Test Here](#)