



Amazon

Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate

NEW QUESTION 1

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 2

After you recommend Amazon Redshift to a client as an alternative solution to paying data warehouses to analyze his data, your client asks you to explain why you are recommending Redshift. Which of the following would be a reasonable response to his request?

- A. It has high performance at scale as data and query complexity grows.
- B. It prevents reporting and analytic processing from interfering with the performance of OLTP workloads.
- C. You don't have the administrative burden of running your own data warehouse and dealing with setup, durability, monitoring, scaling, and patching.
- D. All answers listed are a reasonable response to his QUESTION

Answer: D

Explanation:

Amazon Redshift delivers fast query performance by using columnar storage technology to improve I/O efficiency and parallelizing queries across multiple nodes. Redshift uses standard PostgreSQL JDBC and ODBC drivers, allowing you to use a wide range of familiar SQL clients. Data load speed scales linearly with cluster size, with integrations to Amazon S3, Amazon DynamoDB, Amazon Elastic MapReduce, Amazon Kinesis or any SSH-enabled host.

AWS recommends Amazon Redshift for customers who have a combination of needs, such as: High performance at scale as data and query complexity grows

Desire to prevent reporting and analytic processing from interfering with the performance of OLTP workloads

Large volumes of structured data to persist and query using standard SQL and existing BI tools Desire to the administrative burden of running one's own data warehouse and dealing with setup, durability, monitoring, scaling and patching

Reference: https://aws.amazon.com/running_databases/#redshift_anchor

NEW QUESTION 3

A user is launching an EC2 instance in the US East region. Which of the below mentioned options is recommended by AWS with respect to the selection of the availability zone?

- A. Always select the AZ while launching an instance
- B. Always select the US-East-1-a zone for HA
- C. Do not select the AZ; instead let AWS select the AZ
- D. The user can never select the availability zone while launching an instance

Answer: C

Explanation:

When launching an instance with EC2, AWS recommends not to select the availability zone (AZ). AWS specifies that the default Availability Zone should be accepted. This is because it enables AWS to select the best Availability Zone based on the system health and available capacity. If the user launches additional instances, only then an Availability Zone should be specified. This is to specify the same or different AZ from the running instances.

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

NEW QUESTION 4

A user is storing a large number of objects on AWS S3. The user wants to implement the search functionality among the objects. How can the user achieve this?

- A. Use the indexing feature of S3.
- B. Tag the objects with the metadata to search on that.
- C. Use the query functionality of S3.
- D. Make your own DB system which stores the S3 metadata for the search functionalit

Answer: D

Explanation:

In Amazon Web Services, AWS S3 does not provide any query facility. To retrieve a specific object the user needs to know the exact bucket / object key. In this case it is recommended to have an own DB system which manages the S3 metadata and key mapping.

Reference: http://media.amazonwebservices.com/AWS_Storage_Options.pdf

NEW QUESTION 5

After setting up a Virtual Private Cloud (VPC) network, a more experienced cloud engineer suggests that to achieve low network latency and high network throughput you should look into setting up a placement group. You know nothing about this, but begin to do some research about it and are especially curious about its limitations. Which of the below statements is wrong in describing the limitations of a placement group?

- A. Although launching multiple instance types into a placement group is possible, this reduces the likelihood that the required capacity will be available for your launch to succeed.
- B. A placement group can span multiple Availability Zones.
- C. You can't move an existing instance into a placement group.

D. A placement group can span peered VPCs

Answer: B

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 Gbps network. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. To provide the lowest latency, and the highest packet-per-second network performance for your placement group, choose an instance type that supports enhanced networking. Placement groups have the following limitations:

The name you specify for a placement group a name must be unique within your AWS account. A placement group can't span multiple Availability Zones. Although launching multiple instance types into a placement group is possible, this reduces the likelihood that the required capacity will be available for your launch to succeed. We recommend using the same instance type for all instances in a placement group.

You can't merge placement groups. Instead, you must terminate the instances in one placement group, and then relaunch those instances into the other placement group.

A placement group can span peered VPCs; however, you will not get full-bisection bandwidth between instances in peered VPCs. For more information about VPC peering connections, see VPC Peering in the Amazon VPC User Guide.

You can't move an existing instance into a placement group. You can create an AM from your existing instance, and then launch a new instance from the AMI into a placement group.

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

NEW QUESTION 6

What is a placement group in Amazon EC2?

- A. It is a group of EC2 instances within a single Availability Zone.
- B. It 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 Zone

Answer: A

Explanation:

A placement group is a logical grouping of instances within a single Availability Zone. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

NEW QUESTION 7

An edge location refers to which Amazon Web Service?

- A. An edge location is refered 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 sever — CloudFront serves the cached copy when available and only makes a request to your origin server if the edge location receMng the browser's request does not have a copy of the file.

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

NEW QUESTION 8

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

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

Answer: D

Explanation:

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

NEW QUESTION 9

Does DynamoDB support in-place atomic updates?

- A. Yes
- B. No
- C. It does support in-place non-atomic updates
- D. It is not defined

Answer: A

Explanation:

DynamoDB supports in-place atomic updates.

Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/WorkingWithItems.html#WorkingWithItems.AtomicCounters>

NEW QUESTION 10

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.

Reference: http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf

NEW QUESTION 10

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

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

Answer: C

Explanation:

AWS Import/Export supports: Import to Amazon S3

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

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

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

NEW QUESTION 15

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

Reference:

<http://aws.amazon.com/about-aws/whats-new/2013/05/30/amazon-route-53-adds-elb-integration-for-dns-failover/>

NEW QUESTION 20

An Elastic IP address (EIP) is a static IP address designed for dynamic cloud computing. With an EIP, you can mask the failure of an instance or software by rapidly remapping the address to another instance in your account. Your EIP is associated with your AWS account, not a particular EC2 instance, and it remains associated with your account until you choose to explicitly release it. By default how many EIPs is each AWS account limited to on a per region basis?

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

Answer: B

Explanation:

By default, all AWS accounts are limited to 5 Elastic IP addresses per region for each AWS account, because public (IPv4) Internet addresses are a scarce public resource. AWS strongly encourages you to use an EIP primarily for load balancing use cases, and use DNS hostnames for all other inter-node communication.

If you feel your architecture warrants additional EIPs, you would need to complete the Amazon EC2 Elastic IP Address Request Form and give reasons as to your need for additional addresses. Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html#using-instance-addressing-limit>

NEW QUESTION 24

In Amazon EC2, partial instance-hours are billed .

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

Answer: D

Explanation:

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

NEW QUESTION 25

You are setting up a VPC and you need to set up a public subnet within that VPC. Which following requirement must be met for this subnet to be considered a public subnet?

- A. Subnet's traffic is not routed to an internet gateway but has its traffic routed to a virtual private gateway.
- B. Subnet's traffic is routed to an internet gateway.
- C. Subnet's traffic is not routed to an internet gateway.
- D. None of these answers can be considered a public subnet

Answer: B

Explanation:

A virtual private cloud (VPC) is a virtual network dedicated to your AWS account. It is logically isolated from other virtual networks in the AWS cloud. You can launch your AWS resources, such as Amazon EC2 instances, into your VPC. You can configure your VPC: you can select its IP address range, create subnets, and configure route tables, network gateways, and security settings.

A subnet is a range of IP addresses in your VPC. You can launch AWS resources into a subnet that you select. Use a public subnet for resources that must be connected to the internet, and a private subnet for resources that won't be connected to the Internet.

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

If a subnet doesn't have a route to the internet gateway, the subnet is known as a private subnet.

If a subnet doesn't have a route to the internet gateway, but has its traffic routed to a virtual private gateway, the subnet is known as a VPN-only subnet.

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

NEW QUESTION 26

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

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

Answer: C

Explanation:

Like Amazon S3, the Amazon Glacier service provides low-cost, secure, and durable storage. But where S3 is designed for rapid retrieval, Glacier is meant to be used as an archival service for data that is not accessed often, and for which retrieval times of several hours are suitable.

Amazon Glacier automatically encrypts the data using AES-256 and stores it durably in an immutable form. Amazon Glacier is designed to provide average annual durability of 99.999999999% for an archive. It stores each archive in multiple facilities and multiple devices. Unlike traditional systems which can require laborious data verification and manual repair, Glacier performs regular, systematic data integrity checks, and is built to be automatically self-healing.

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

NEW QUESTION 28

Your EBS volumes do not seem to be performing as expected and your team leader has requested you look into improving their performance. Which of the following is not a true statement relating to the performance of your EBS volumes?

- A. Frequent snapshots provide a higher level of data durability and they will not degrade the performance of your application while the snapshot is in progress.
- B. General Purpose (SSD) and Provisioned IOPS (SSD) volumes have a throughput limit of 128 MB/s per volume.
- C. There is a relationship between the maximum performance of your EBS volumes, the amount of I/O you are doing to them, and the amount of time it takes for each transaction to complete.
- D. There is a 5 to 50 percent reduction in IOPS when you first access each block of data on a newly created or restored EBS volume

Answer: A

Explanation:

Several factors can affect the performance of Amazon EBS volumes, such as instance configuration, I/O characteristics, workload demand, and storage configuration.

Frequent snapshots provide a higher level of data durability, but they may slightly degrade the

performance of your application while the snapshot is in progress. This trade off becomes critical when you have data that changes rapidly. Whenever possible, plan for snapshots to occur during off-peak times in order to minimize workload impact.

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

NEW QUESTION 31

You've created your first load balancer and have registered your EC2 instances with the load balancer. Elastic Load Balancing routinely performs health checks on all the registered EC2 instances and automatically distributes all incoming requests to the DNS name of your load balancer across your registered, healthy EC2 instances. By default, the load balancer uses the _ protocol for checking the health of your instances.

- A. HTTPS
- B. HTTP
- C. ICMP
- D. IPv6

Answer: B

Explanation:

In Elastic Load Balancing a health configuration uses information such as protocol, ping port, ping path (URL), response timeout period, and health check interval to determine the health state of the instances registered with the load balancer.

Currently, HTTP on port 80 is the default health check. Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/TerminologyandKeyConcepts.html>

NEW QUESTION 32

is a fast, filexible, fully managed push messaging service.

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

Answer: A

Explanation:

Amazon Simple Notification Service (Amazon SNS) is a fast, filexible, fully managed push messaging service. Amazon SNS makes it simple and cost-effective to push to mobile devices such as iPhone, iPad, Android, Kindle Fire, and internet connected smart devices, as well as pushing to other distributed services.

Reference: http://aws.amazon.com/sns/?nc1=h_l2_as

NEW QUESTION 36

In an experiment, if the minimum size for an Auto Scaling group is 1 instance, which of the following statements holds true when you terminate the running instance?

- A. Auto Scaling must launch a new instance to replace it.
- B. Auto Scaling will raise an alarm and send a notification to the user for action.
- C. Auto Scaling must configure the schedule actMty that terminates the instance after 5 days.
- D. Auto Scaling will terminate the experimen

Answer: A

Explanation:

If the minimum size for an Auto Scaling group is 1 instance, when you terminate the running instance, Auto Scaling must launch a new instance to replace it.

Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html

NEW QUESTION 40

An organization has created an application which is hosted on the AWS EC2 instance. The application stores images to S3 when the end user uploads to it. The organization does not want to store the AWS secure credentials required to access the S3 inside the instance. Which of the below mentioned options is a possible solution to avoid any security threat?

- A. Use the IAM based single sign between the AWS resources and the organization application.
- B. Use the IAM role and assign it to the instance.
- C. Since the application is hosted on EC2, it does not need credentials to access S3.
- D. Use the X.509 certificates instead of the access and the secret access key

Answer: B

Explanation:

The AWS IAM role uses temporary security credentials to access AWS services. Once the role is assigned to an instance, it will not need any security credentials to be stored on the instance. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

NEW QUESTION 44

You are running PostgreSQL on Amazon RDS and it seems to be all running smoothly deployed in one availability zone. A database administrator asks you if DB instances running PostgreSQL support Multi-AZ deployments. What would be a correct response to this QUESTION ?

- A. Yes.
- B. Yes but only for small db instances.
- C. No.
- D. Yes but you need to request the service from AW

Answer: A

Explanation:

Amazon RDS supports DB instances running several versions of PostgreSQL. Currently we support PostgreSQL versions 9.3.1, 9.3.2, and 9.3.3. You can create DB instances and DB snapshots, point-in-time restores and backups.

DB instances running PostgreSQL support Multi-AZ deployments, Provisioned IOPS, and can be created inside a VPC. You can also use SSL to connect to a DB instance running PostgreSQL.

You can use any standard SQL client application to run commands for the instance from your client computer. Such applications include pgAdmin, a popular Open Source administration and development tool for PostgreSQL, or psql, a command line utility that is part of a PostgreSQL installation. In order to deliver a managed service experience, Amazon RDS does not provide host access to DB instances, and it restricts access to certain system procedures and tables that require advanced privileges. Amazon RDS supports access to databases on a DB instance using any standard SQL client application. Amazon RDS does not allow direct host access to a DB instance via Telnet or Secure Shell (SSH).

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

NEW QUESTION 49

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

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

Answer: A

Explanation:

A placement group is a logical grouping of EC2 instances within a single Availability Zone. Using placement groups enables applications to participate in a low-latency, 10 Gbps network. Placement groups are recommended for applications that benefit from low network latency, high network throughput or both.

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

NEW QUESTION 53

An organization has developed a mobile application which allows end users to capture a photo on their mobile device, and store it inside an application. The application internally uploads the data to AWS S3. The organization wants each user to be able to directly upload data to S3 using their Google ID. How will the mobile app allow this?

- A. Use the AWS Web identity federation for mobile applications, and use it to generate temporary security credentials for each user.
- B. It is not possible to connect to AWS S3 with a Google ID.
- C. Create an IAM user every time a user registers with their Google ID and use IAM to upload files to S3.
- D. Create a bucket policy with a condition which allows everyone to upload if the login ID has a Google part to it.

Answer: A

Explanation:

For Amazon Web Services, the Web identity federation allows you to create cloud-backed mobile apps that use public identity providers, such as login with Facebook, Google, or Amazon. It will create temporary security credentials for each user, which will be authenticated by the AWS services, such as S3.

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

NEW QUESTION 56

You have been doing a lot of testing of your VPC Network by deliberately failing EC2 instances to test whether instances are failing over properly. Your customer who will be paying the AWS bill for all this asks you if he being charged for all these instances. You try to explain to him how the billing works on EC2 instances to the best of your knowledge. What would be an appropriate response to give to the customer in regards to this?

- A. Billing commences when Amazon EC2 AM instance is completely up and billing ends as soon as the instance starts to shutdown.
- B. Billing only commences only after 1 hour of uptime and billing ends when the instance terminates.
- C. Billing commences when Amazon EC2 initiates the boot sequence of an AM instance and billing ends when the instance shuts down.
- D. Billing commences when Amazon EC2 initiates the boot sequence of an AM instance and billing ends as soon as the instance starts to shutdown.

Answer: C

Explanation:

Billing commences when Amazon EC2 initiates the boot sequence of an AM instance. Billing ends when the instance shuts down, which could occur through a web services command, by running "shutdown -h", or through instance failure.

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

NEW QUESTION 59

Once again your customers are concerned about the security of their sensitive data and with their latest enquiry ask about what happens to old storage devices on AWS. What would be the best answer to this QUESTION ?

- A. AWS reformats the disks and uses them again.
- B. AWS uses the techniques detailed in DoD 5220.22-M to destroy data as part of the decommissioning process.
- C. AWS uses their own proprietary software to destroy data as part of the decommissioning process.
- D. AWS uses a 3rd party security organization to destroy data as part of the decommissioning process

Answer: B

Explanation:

When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals.

AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual ") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process.

All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.
Reference: <http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf>

NEW QUESTION 63

Your company has been storing a lot of data in Amazon Glacier and has asked for an inventory of what is in there exactly. So you have decided that you need to download a vault inventory. Which of the following statements is incorrect in relation to Vault Operations in Amazon Glacier?

- A. You can use Amazon Simple Notification Service (Amazon SNS) notifications to notify you when the job completes.
- B. A vault inventory refers to the list of archives in a vault.
- C. You can use Amazon Simple Queue Service (Amazon SQS) notifications to notify you when the job completes.
- D. Downloading a vault inventory is an asynchronous operation

Answer: C

Explanation:

Amazon Glacier supports various vault operations.

A vault inventory refers to the list of archives in a vault. For each archive in the list, the inventory provides archive information such as archive ID, creation date, and size. Amazon Glacier updates the vault inventory approximately once a day, starting on the day the first archive is uploaded to the vault. A vault inventory must exist for you to be able to download it.

Downloading a vault inventory is an asynchronous operation. You must first initiate a job to download the inventory. After receiving the job request, Amazon Glacier prepares your inventory for download. After the job completes, you can download the inventory data.

Given the asynchronous nature of the job, you can use Amazon Simple Notification Service (Amazon SNS) notifications to notify you when the job completes. You can specify an Amazon SNS topic for each individual job request or configure your vault to send a notification when specific vault events occur. Amazon Glacier prepares an inventory for each vault periodically, every 24 hours. If there have been no archive additions or deletions to the vault since the last inventory, the inventory date is not updated. When you initiate a job for a vault inventory, Amazon Glacier returns the last inventory it generated, which is a point-in-time snapshot and not real-time data. You might not find it useful to retrieve vault inventory for each archive upload. However, suppose you maintain a database on the client-side associating metadata about the archives you upload to Amazon Glacier. Then, you might find the vault inventory useful to reconcile information in your database with the actual vault inventory.

Reference: <http://docs.aws.amazon.com/amazonglacier/latest/dev/working-with-vaults.html>

NEW QUESTION 66

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

NEW QUESTION 69

You need to change some settings on Amazon Relational Database Service but you do not want the database to reboot immediately which you know might happen depending on the setting that you change. Which of the following will cause an immediate DB instance reboot to occur?

- A. You change storage type from standard to PIOPS, and Apply Immediately is set to true.
- B. You change the DB instance class, and Apply Immediately is set to false.
- C. You change a static parameter in a DB parameter group.
- D. You change the backup retention period for a DB instance from 0 to a nonzero value or from a nonzero value to 0, and Apply Immediately is set to false.

Answer: A

Explanation:

A DB instance outage can occur when a DB instance is rebooted, when the DB instance is put into a state that prevents access to it, and when the database is restarted. A reboot can occur when you manually reboot your DB instance or when you change a DB instance setting that requires a reboot before it can take effect.

A DB instance reboot occurs immediately when one of the following occurs:

You change the backup retention period for a DB instance from 0 to a nonzero value or from a nonzero value to 0 and set Apply Immediately to true.

You change the DB instance class, and Apply Immediately is set to true.

You change storage type from standard to PIOPS, and Apply Immediately is set to true.

A DB instance reboot occurs during the maintenance window when one of the following occurs:

You change the backup retention period for a DB instance from 0 to a nonzero value or from a nonzero value to 0, and Apply Immediately is set to false.

You change the DB instance class, and Apply Immediately is set to false. Reference:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Troubleshooting.Security

NEW QUESTION 74

What does the following policy for Amazon EC2 do?

```
{
  "Statement": [{
    "Effect": "Allow", "Action": "ec2:Describe*", "Resource": "*"
  }]
}
```



```
}
```

- A. Allow users to use actions that start with "Describe" over all the EC2 resources.
- B. Share an AMI with a partner
- C. Share an AMI within the account
- D. Allow a group to only be able to describe, run, stop, start, and terminate instances

Answer: A

Explanation:

You can use IAM policies to control the actions that your users can perform against your EC2 resources. For instance, a policy with the following statement will allow users to perform actions whose name start with "Describe" against all your EC2 resources.

```
{  
  "Statement": [{  
    "Effect": "Allow", "Action": "ec2:Describe*", "Resource": "*" } ]  
}
```

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

NEW QUESTION 77

Amazon RDS provides high availability and failover support for DB instances using .

- A. customized deployments
- B. Appstream customizations
- C. log events
- D. Multi-AZ deployments

Answer: D

Explanation:

Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. Multi-AZ deployments for Oracle, PostgreSQL, MySQL, and MariaDB DB instances use Amazon technology, while SQL Server DB instances use SQL Server Mirroring.

Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.IV|u|tiAZ.html>

NEW QUESTION 79

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.

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

NEW QUESTION 83

What does Amazon DynamoDB provide?

- A. A predictable and scalable MySQL database
- B. A fast and reliable PL/SQL database cluster
- C. A standalone Cassandra database, managed by Amazon Web Services
- D. A fast, highly scalable managed NoSQL database service

Answer: D

Explanation:

Amazon DynamoDB is a managed NoSQL database service offered by Amazon. It automatically manages tasks like scalability for you while it provides high availability and durability for your data, allowing you to concentrate in other aspects of your application.

Reference: check link - https://aws.amazon.com/running_databases/

NEW QUESTION 88

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).

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

NEW QUESTION 90

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-Classic 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-Classic, 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

Reference: http://media.amazonwebservices.com/AWS_Cloud_Best_Practices.pdf

NEW QUESTION 93

Amazon S3 allows you to set per-file permissions to grant read and/or write access. However you have decided that you want an entire bucket with 100 files already in it to be accessible to the public. You don't want to go through 100 files individually and set permissions. What would be the best way to do this?

- A. Move the bucket to a new region
- B. Add a bucket policy to the bucket.
- C. Move the files to a new bucket.
- D. Use Amazon EBS instead of S3

Answer: B

Explanation:

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

Amazon S3 provides four different access control mechanisms: AWS Identity and Access Management (IAM) policies, Access Control Lists (ACLs), bucket policies, and query string authentication. IAM enables organizations to create and manage multiple users under a single AWS account. With IAM policies, you can grant IAM users fine-grained control to your Amazon S3 bucket or objects. You can use ACLs to selectively add (grant) certain permissions on individual objects.

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

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

Reference: <http://aws.amazon.com/s3/details/#security>

NEW QUESTION 98

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 101

You have been using T2 instances as your CPU requirements have not been that intensive. However you now start to think about larger instance types and start looking at M and IV3 instances. You are a little confused as to the differences between them as they both seem to have the same ratio of CPU and memory.

Which statement below is incorrect as to why you would use one over the other?

- A. M3 instances are less expensive than M1 instances.
- B. IV|3 instances are configured with more swap memory than M instances.
- C. IV|3 instances provide better, more consistent performance than M instances for most use-cases.
- D. M3 instances also offer SSD-based instance storage that delivers higher I/O performance.

Answer: B

Explanation:

Amazon EC2 allows you to set up and configure everything about your instances from your operating system up to your applications. An Amazon Machine Image (AMI) is simply a packaged-up environment that includes all the necessary bits to set up and boot your instance. M1 and M3 Standard instances have the same ratio of CPU and memory, some reasons below as to why you would use one over the other. IV|3 instances provide better, more consistent performance than M instances for most use-cases. M3 instances also offer SSD-based instance storage that delivers higher I/O performance. M3 instances are also less expensive than M1 instances. Due to these reasons, we recommend M3 for applications that require general purpose instances with a balance of compute, memory, and network resources. However, if you need more disk storage than what is provided in M3 instances, you may still find M1 instances useful for running your applications. Reference: <https://aws.amazon.com/ec2/faqs/>

NEW QUESTION 102

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. Reference: <http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/TerminologyandKeyConcepts.html#session-stickiness>

NEW QUESTION 105

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

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

Answer: C

Explanation:

Amazon Relational Database Service (Amazon RDS) is a web service that makes it easier to set up, operate, and scale a relational database in the cloud. It provides cost-efficient, resizable capacity for an industry-standard relational database and manages common database administration tasks. There are several ways you can track the performance and health of a database or a DB instance. You can: Use the free Amazon CloudWatch service to monitor the performance and health of a DB instance. Subscribe to Amazon RDS events to be notified when changes occur with a DB instance, DB snapshot, DB parameter group, or DB security group. View, download, or watch database log files using the Amazon RDS console or Amazon RDS APIs. You can also query some database log files that are loaded into database tables. Use the AWS CloudTrail service to record AWS calls made by your AWS account. The calls are recorded in log files and stored in an Amazon S3 bucket. Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Monitoring.html

NEW QUESTION 107

You are building a system to distribute confidential documents to employees. Using CloudFront, what method could be used to serve content that is stored in S3, but not publically accessible from S3 directly?

- A. Add the CloudFront account security group "amazon-cf/amazon-cf-sg" to the appropriate S3 bucket policy.
- B. Create a S3 bucket policy that lists the CloudFront distribution ID as the Principal and the target bucket as the Amazon Resource Name (ARN).
- C. Create an Identity and Access Management (IAM) User for CloudFront and grant access to the objects in your S3 bucket to that IAM User.
- D. Create an Origin Access Identity (OAI) for CloudFront and grant access to the objects in your S3 bucket to that OAI.

Answer: D

Explanation:

You restrict access to Amazon S3 content by creating an origin access identity, which is a special CloudFront user. You change Amazon S3 permissions to give the origin access identity permission to access your objects, and to remove permissions from everyone else. When your users access your Amazon S3 objects using CloudFront URLs, the CloudFront origin access identity gets the objects on your users' behalf. If your users try to access objects using Amazon S3 URLs,

they're denied access. The origin access identity has permission to access objects in your Amazon S3 bucket, but users don't. Reference: <http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html>

NEW QUESTION 112

A user is aware that a huge download is occurring on his instance. He has already set the Auto Scaling policy to increase the instance count when the network I/O increases beyond a certain limit. How can the user ensure that this temporary event does not result in scaling?

- A. The network I/O are not affected during data download
- B. The policy cannot be set on the network I/O
- C. There is no way the user can stop scaling as it is already configured
- D. Suspend scaling

Answer: D

Explanation:

The user may want to stop the automated scaling processes on the Auto Scaling groups either to perform manual operations or during emergency situations. To perform this, the user can suspend one or more scaling processes at any time. Once it is completed, the user can resume all the suspended processes. Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html

NEW QUESTION 114

An accountant asks you to design a small VPC network for him and, due to the nature of his business, just needs something where the workload on the network will be low, and dynamic data will be accessed infrequently. Being an accountant, low cost is also a major factor. Which EBS volume type would best suit his requirements?

- A. Magnetic
- B. Any, as they all perform the same and cost the same.
- C. General Purpose (SSD)
- D. Magnetic or Provisioned IOPS (SSD)

Answer: A

Explanation:

You can choose between three EBS volume types to best meet the needs of their workloads: General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic. General Purpose (SSD) is the new, SSD-backed, general purpose EBS volume type that we recommend as the default choice for customers. General Purpose (SSD) volumes are suitable for a broad range of workloads, including small to medium sized databases, development and test environments, and boot volumes. Provisioned IOPS (SSD) volumes offer storage with consistent and low-latency performance, and are designed for I/O intensive applications such as large relational or NoSQL databases. Magnetic volumes provide the lowest cost per gigabyte of all EBS volume types. Magnetic volumes are ideal for workloads where data is accessed infrequently, and applications where the lowest storage cost is important.

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

NEW QUESTION 118

A user has launched one EC2 instance in the US East region and one in the US West region. The user has launched an RDS instance in the US East region. How can the user configure access from both the EC2 instances to RDS?

- A. It is not possible to access RDS of the US East region from the US West region
- B. Configure the US West region's security group to allow a request from the US East region's instance and configure the RDS security group's ingress rule for the US East EC2 group
- C. Configure the security group of the US East region to allow traffic from the US West region's instance and configure the RDS security group's ingress rule for the US East EC2 group
- D. Configure the security group of both instances in the ingress rule of the RDS security group

Answer: C

Explanation:

The user cannot authorize an Amazon EC2 security group if it is in a different AWS Region than the RDS DB instance. The user can authorize an IP range or specify an Amazon EC2 security group in the same region that refers to an IP address in another region. In this case allow IP of US West inside US East's security group and open the RDS security group for US East region.

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

NEW QUESTION 122

A client of yours has a huge amount of data stored on Amazon S3, but is concerned about someone stealing it while it is in transit. You know that all data is encrypted in transit on AWS, but which of the following is wrong when describing server-side encryption on AWS?

- A. Amazon S3 server-side encryption employs strong multi-factor encryption.
- B. Amazon S3 server-side encryption uses one of the strongest block ciphers available, 256-bit Advanced Encryption Standard (AES-256), to encrypt your data.
- C. In server-side encryption, you manage encryption/decryption of your data, the encryption keys, and related tools.
- D. Server-side encryption is about data encryption at rest—that is, Amazon S3 encrypts your data as it writes it to disks.

Answer: C

Explanation:

Amazon S3 encrypts your object before saving it on disks in its data centers and decrypts it when you download the objects. You have two options depending on how you choose to manage the encryption keys: Server-side encryption and client-side encryption.

Server-side encryption is about data encryption at rest—that is, Amazon S3 encrypts your data as it writes it to disks in its data centers and decrypts it for you when you access it. As long as you authenticate your request and you have access permissions, there is no difference in the way you access encrypted or unencrypted objects. Amazon S3 manages encryption and decryption for you. For example, if you share your objects using a pre-signed URL, that URL works the same way for both encrypted and unencrypted objects.

In client-side encryption, you manage encryption/decryption of your data, the encryption keys, and related tools. Server-side encryption is an alternative to client-side encryption in which Amazon S3 manages the encryption of your data, freeing you from the tasks of managing encryption and encryption keys.

Amazon S3 server-side encryption 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.

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

NEW QUESTION 125

You have just set up a large site for a client which involved a huge database which you set up with Amazon RDS to run as a Multi-AZ deployment. You now start to worry about what will happen if the database instance fails. Which statement best describes how this database will function if there is a database failure?

- A. Updates to your DB Instance are synchronously replicated across Availability Zones to the standby in order to keep both in sync and protect your latest database updates against DB Instance failure.
- B. Your database will not resume operation without manual administrative intervention.
- C. Updates to your DB Instance are asynchronously replicated across Availability Zones to the standby in order to keep both in sync and protect your latest database updates against DB Instance failure.
- D. Updates to your DB Instance are synchronously replicated across S3 to the standby in order to keep both in sync and protect your latest database updates against DB Instance failure.

Answer: A

Explanation:

Amazon Relational Database Service (Amazon RDS) is a managed service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity, while managing time-consuming database administration tasks, freeing you up to focus on your applications and business.

When you create or modify your DB Instance to run as a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous "standby" replica in a different Availability Zone. Updates to your DB Instance are synchronously replicated across Availability Zones to the standby in order to keep both in sync and protect your latest database updates against DB Instance failure.

During certain types of planned maintenance, or in the unlikely event of DB Instance failure or Availability Zone failure, Amazon RDS will automatically failover to the standby so that you can resume database writes and reads as soon as the standby is promoted. Since the name record for your DB Instance remains the same, your application can resume database operation without the need for manual administrative intervention. With Multi-AZ deployments, replication is transparent: you do not interact directly with the standby, and it cannot be used to serve read traffic. If you are using Amazon RDS for MySQL and are looking to scale read traffic beyond the capacity constraints of a single DB Instance, you can deploy one or more Read Replicas.

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

NEW QUESTION 126

You are signed in as root user on your account but there is an Amazon S3 bucket under your account that you cannot access. What is a possible reason for this?

- A. An IAM user assigned a bucket policy to an Amazon S3 bucket and didn't specify the root user as a principal
- B. The S3 bucket is full.
- C. The S3 bucket has reached the maximum number of objects allowed.
- D. You are in the wrong availability zone

Answer: A

Explanation:

With IAM, you can centrally manage users, security credentials such as access keys, and permissions that control which AWS resources users can access.

In some cases, you might have an IAM user with full access to IAM and Amazon S3. If the IAM user assigns a bucket policy to an Amazon S3 bucket and doesn't specify the root user as a principal, the root user is denied access to that bucket. However, as the root user, you can still access the bucket by modifying the bucket policy to allow root user access.

Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/iam-troubleshooting.html#testing2>

NEW QUESTION 128

A scope has been handed to you to set up a super fast gaming server and you decide that you will use Amazon DynamoDB as your database. For efficient access to data in a table, Amazon DynamoDB creates and maintains indexes for the primary key attributes. A secondary index is a data structure that contains a subset of attributes from a table, along with an alternate key to support Query operations. How many types of secondary indexes does DynamoDB support?

- A. 2
- B. 16
- C. 4
- D. As many as you need

Answer: A

Explanation:

DynamoDB supports two types of secondary indexes:

Local secondary index — an index that has the same hash key as the table, but a different range key. A local secondary index is "local" in the sense that every partition of a local secondary index is scoped to a table partition that has the same hash key.

Global secondary index — an index with a hash and range key that can be different from those on the table. A global secondary index is considered "global" because queries on the index can span all of the data in a table, across all partitions.

Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html>

NEW QUESTION 130

You need to set up security for your VPC and you know that Amazon VPC provides two features that you can use to increase security for your VPC: security groups and network access control lists (ACLs). You have already looked into security groups and you are now trying to understand ACLs. Which statement below is incorrect in relation to ACLs?

- A. Supports allow rules and deny rules.
- B. Is stateful: Return traffic is automatically allowed, regardless of any rules.
- C. Processes rules in number order when deciding whether to allow traffic.
- D. Operates at the subnet level (second layer of defense).

Answer: B

Explanation:

Amazon VPC provides two features that you can use to increase security for your VPC:

Security groups—Act as a firewall for associated Amazon EC2 instances, controlling both inbound and outbound traffic at the instance level

Network access control lists (ACLs)—Act as a firewall for associated subnets, controlling both inbound and outbound traffic at the subnet level

Security groups are stateful: (Return traffic is automatically allowed, regardless of any rules) Network ACLs are stateless: (Return traffic must be explicitly allowed by rules)

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

NEW QUESTION 133

While creating an Amazon RDS DB, your first task is to set up a DB that controls which IP address or EC2 instance can access your DB Instance.

- A. security token pool
- B. security token
- C. security pool
- D. security group

Answer: D

Explanation:

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

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

NEW QUESTION 136

You are very concerned about security on your network because you have multiple programmers testing APIs and SDKs and you have no idea what is happening. You think CloudTrail may help but are not sure what it does. Which of the following statements best describes the AWS service CloudTrail?

- A. With AWS CloudTrail you can get a history of AWS API calls and related events for your account.
- B. With AWS CloudTrail you can get a history of IAM users for your account.
- C. With AWS CloudTrail you can get a history of S3 logfiles for your account.
- D. With AWS CloudTrail you can get a history of CloudFormation JSON scripts used for your account.

Answer: A

Explanation:

With AWS CloudTrail, you can get a history of AWS API calls for your account, including API calls made via the AWS Management Console, the AWS SDKs, the command line tools, and higher-level AWS services. You can also identify which users and accounts called AWS APIs for services that support CloudTrail, the source IP address the calls were made from, and when the calls occurred.

You can identify which users and accounts called AWS for services that support CloudTrail, the source IP address the calls were made from, and when the calls occurred. You can integrate CloudTrail into applications using the API, automate trail creation for your organization, check the status of your trails, and control how administrators turn CloudTrail logging on and off.

Reference: http://docs.aws.amazon.com/awscloudtrail/latest/userguide/what_is_cloud_trail_top_level.html

NEW QUESTION 137

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 142

You receive a bill from AWS but are confused because you see you are incurring different costs for the exact same storage size in different regions on Amazon S3. You ask AWS why this is so. What response would you expect to receive from AWS?

- A. We charge less in different time zones.
- B. We charge less where our costs are less.
- C. This will balance out next bill.
- D. It must be a mistake.

Answer: B

Explanation:

Amazon S3 is storage for the internet. It's a simple storage service that offers software developers a highly-scalable, reliable, and low-latency data storage infrastructure at very low costs.

AWS charges less where their costs are less.

For example, their costs are lower in the US Standard Region than in the US West (Northern California) Region.
Reference: <https://aws.amazon.com/s3/faqs/>

NEW QUESTION 145

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

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

Answer: B

Explanation:

In Amazon VPC, you can assign any Private IP address to your instance as long as it is: Part of the associated subnet's IP address range
Not reserved by Amazon for IP networking purposes Not currently assigned to another interface
Reference: <http://aws.amazon.com/vpc/faqs/>

NEW QUESTION 148

You are planning and configuring some EBS volumes for an application. In order to get the most performance out of your EBS volumes, you should attach them to an instance with enough to support your volumes.

- A. Redundancy
- B. Storage
- C. Bandwidth
- D. Memory

Answer: C

Explanation:

When you plan and configure EBS volumes for your application, it is important to consider the configuration of the instances that you will attach the volumes to. In order to get the most performance out of your EBS volumes, you should attach them to an instance with enough bandwidth to support your volumes, such as an EBS-optimized instance or an instance with 10 Gigabit network connectMty. This is especially important when you use General Purpose (SSD) or Provisioned IOPS (SSD) volumes, or when you stripe multiple volumes together in a RAID configuration.
Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-ec2-config.html>

NEW QUESTION 151

Can a single EBS volume be attached to multiple EC2 instances at the same time?

- A. Yes
- B. No
- C. Only for high-performance EBS volumes.
- D. Only when the instances are located in the US region

Answer: B

Explanation:

You can't attach an EBS volume to multiple EC2 instances. This is because it is equivalent to using a single hard drive with many computers at the same time.
Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html>

NEW QUESTION 153

You have been asked to tighten up the password policies in your organization after a serious security breach, so you need to consider every possible security measure. Which of the following is not an account password policy for IAM Users that can be set?

- A. Force IAM users to contact an account administrator when the user has allowed his or her password to expue.
- B. A minimum password length.
- C. Force IAM users to contact an account administrator when the user has entered his password incorrectly.
- D. Prevent IAM users from reusing previous password

Answer: C

Explanation:

IAM users need passwords in order to access the AWS Management Console. (They do not need passwords if they will access AWS resources programmatically by using the CLI, AWS SDKs, or the APIs.)
You can use a password policy to do these things: Set a minimum password length.
Require specific character types, including uppercase letters, lowercase letters, numbers, and non-alphanumeric characters. Be sure to remind your users that passwords are case sensitive. Allow all IAM users to change their own passwords.
Require IAM users to change their password after a specified period of time (enable password expiration). Prevent IAM users from reusing previous passwords.
Force IAM users to contact an account administrator when the user has allowed his or her password to expue.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_ManagingPasswordPolicies.html

NEW QUESTION 155

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

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

D. It will warn the user about the error and ask the user to manually create RDS

Answer: A

Explanation:

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

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

NEW QUESTION 157

You have just finished setting up an advertisement server in which one of the obvious choices for a service was Amazon Elastic Map Reduce (EMR) and are now troubleshooting some weird cluster states that you are seeing. Which of the below is not an Amazon EMR cluster state?

- A. STARTING
- B. STOPPED
- C. RUNNING
- D. WAITING

Answer: B

Explanation:

Amazon Elastic Map Reduce (EMR) is a web service that enables businesses, researchers, data analysts, and developers to easily and cost-effectively process vast amounts of data.

Amazon EMR historically referred to an Amazon EMR cluster (and all processing steps assigned to it) as a "cluster". Every cluster has a unique identifier that starts with "j-".

The different cluster states of an Amazon EMR cluster are listed below. STARTING — The cluster provisions, starts, and configures EC2 instances.

BOOTSTRAPPING — Bootstrap actions are being executed on the cluster. RUNNING — A step for the cluster is currently being run.

WAITING — The cluster is currently active, but has no steps to run. TERMINATING - The cluster is in the process of shutting down. TERMINATED - The cluster was shut down without error. TERMINATED_WITH_ERRORS - The cluster was shut down with errors.

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

NEW QUESTION 162

Is it possible to get a history of all EC2 API calls made on your account for security analysis and operational troubleshooting purposes?

- A. Yes, by default, the history of your API calls is logged.
- B. Yes, you should turn on the CloudTrail in the AWS console.
- C. No, you can only get a history of VPC API calls.
- D. No, you cannot store history of EC2 API calls on Amazon.

Answer: B

Explanation:

To get a history of all EC2 API calls (including VPC and EBS) made on your account, you simply turn on CloudTrail in the AWS Management Console.

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

NEW QUESTION 164

You have just set up your first Elastic Load Balancer (ELB) but it does not seem to be configured properly. You discover that before you start using ELB, you have to configure the listeners for your load balancer. Which protocols does ELB use to support the load balancing of applications?

- A. HTTP and HTTPS
- B. HTTP, HTTPS, TCP, SSL and SSH
- C. HTTP, HTTPS, TCP, and SSL
- D. HTTP, HTTPS, TCP, SSL and SFTP

Answer: C

Explanation:

Before you start using Elastic Load Balancing (ELB), you have to configure the listeners for your load balancer. A listener is a process that listens for connection requests. It is configured with a protocol and a port number for front-end (client to load balancer) and back-end (load balancer to back-end instance) connections.

Elastic Load Balancing supports the load balancing of applications using HTTP, HTTPS (secure HTTP), TCP, and SSL (secure TCP) protocols. The HTTPS uses the SSL protocol to establish secure connections over the HTTP layer. You can also use SSL protocol to establish secure connections over the TCP layer.

The acceptable ports for both HTTPS/SSL and HTTP/TCP connections are 25, 80, 443, 465, 587, and 1024-65535.

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-listener-config.html>

NEW QUESTION 166

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

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

Answer: C

Explanation:

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

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

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

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

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

NEW QUESTION 170

You have just discovered that you can upload your objects to Amazon S3 using Multipart Upload API. You start to test it out but are unsure of the benefits that it would provide. Which of the following is not a benefit of using multipart uploads?

- A. You can begin an upload before you know the final object size.
- B. Quick recovery from any network issues.
- C. Pause and resume object uploads.
- D. It's more secure than normal upload

Answer: D

Explanation:

Multipart upload in Amazon S3 allows you to upload a single object as a set of parts. Each part is a contiguous portion of the object's data. You can upload these object parts independently and in any order.

If transmission of any part fails, you can re-transmit that part without affecting other parts. After all parts of your object are uploaded, Amazon S3 assembles these parts and creates the object. In general, when your object size reaches 100 MB, you should consider using multipart uploads instead of uploading the object in a single operation.

Using multipart upload provides the following advantages:

Improved throughput—You can upload parts in parallel to improve throughput.

Quick recovery from any network issues—Smaller part size minimizes the impact of restarting a failed upload due to a network error.

Pause and resume object uploads—You can upload object parts over time. Once you initiate a multipart upload there is no expiry; you must explicitly complete or abort the multipart upload.

Begin an upload before you know the final object size—You can upload an object as you are creating it. Reference:

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

NEW QUESTION 173

Which of the following would you use to list your AWS Import/Export jobs?

- A. Amazon RDS
- B. AWS Import/Export Web Service Tool
- C. Amazon S3 REST API
- D. AWS Elastic Beanstalk

Answer: C

Explanation:

You can list AWS Import/Export jobs with the ListJobs command using the command line client or REST API.

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

NEW QUESTION 177

You need to set up security for your VPC and you know that Amazon VPC provides two features that you can use to increase security for your VPC: Security groups and network access control lists (ACLs). You start to look into security groups first. Which statement below is incorrect in relation to security groups?

- A. Are stateful: Return traffic is automatically allowed, regardless of any rules.
- B. Evaluate all rules before deciding whether to allow traffic.
- C. Support allow rules and deny rules.
- D. Operate at the instance level (first layer of defense).

Answer: C

Explanation:

Amazon VPC provides two features that you can use to increase security for your VPC:

Security groups—Act as a firewall for associated Amazon EC2 instances, controlling both inbound and outbound traffic at the instance level and supports allow rules only.

Network access control lists (ACLs)—Act as a firewall for associated subnets, controlling both inbound and outbound traffic at the subnet level and supports allow rules and deny rules.

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

NEW QUESTION 181

You are setting up some IAM user policies and have also become aware that some services support resource-based permissions, which let you attach policies to the service's resources instead of to IAM users or groups. Which of the below statements is true in regards to resource-level permissions?

- A. All services support resource-level permissions for all actions.
- B. Resource-level permissions are supported by Amazon CloudFront
- C. All services support resource-level permissions only for some actions.
- D. Some services support resource-level permissions only for some action

Answer: D

Explanation:

AWS Identity and Access Management is a web service that enables Amazon Web Services (AWS) customers to manage users and user permissions in AWS. The service is targeted at organizations with multiple users or systems that use AWS products such as Amazon EC2, Amazon RDS, and the AWS Management Console. With IAM, you can centrally manage users, security credentials such as access keys, and permissions that control which AWS resources users can access.

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

The resource-level permissions service supports IAM policies in which you can specify individual resources using Amazon Resource Names (ARNs) in the policy's Resource element.

Some services support resource-level permissions only for some actions.

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

NEW QUESTION 182

You are using Amazon SES as an email solution but are unsure of what its limitations are. Which statement below is correct in regards to that?

- A. New Amazon SES users who have received production access can send up to 1,000 emails per 24-hour period, at a maximum rate of 10 emails per second.
- B. Every Amazon SES sender has a the same set of sending limits
- C. Sending limits are based on messages rather than on recipients
- D. Every Amazon SES sender has a unique set of sending limits

Answer: D

Explanation:

Amazon Simple Email Service (Amazon SES) is a highly scalable and cost-effective email-sending service for businesses and developers. Amazon SES eliminates the complexity and expense of building an in-house email solution or licensing, installing, and operating a third-party email service for this type of email communication.

Every Amazon SES sender has a unique set of sending limits, which are calculated by Amazon SES on an ongoing basis:

Sending quota — the maximum number of emails you can send in a 24-hour period. Maximum send rate — the maximum number of emails you can send per second.

New Amazon SES users who have received production access can send up to 10,000 emails per 24-hour period, at a maximum rate of 5 emails per second.

Amazon SES automatically adjusts these limits upward, as long as you send high-quality email. If your existing quota is not adequate for your needs and the system has not automatically increased your quota, you can submit an SES Sending Quota Increase case at any time.

Sending limits are based on recipients rather than on messages. You can check your sending limits at any time by using the Amazon SES console.

Note that if your email is detected to be of poor or QUESTION able quality (e.g., high complaint rates, high bounce rates, spam, or abusive content), Amazon SES might temporarily or permanently reduce your permitted send volume, or take other action as AWS deems appropriate.

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

NEW QUESTION 186

You have set up an S3 bucket with a number of images in it and you have decided that you want anybody to be able to access these images, even anonymous users. To accomplish this you create a bucket policy. You will need to use an Amazon S3 bucket policy that specifies a in the principal element, which means anyone can access the bucket.

- A. hash tag (#)
- B. anonymous user
- C. wildcard (*)
- D. S3 user

Answer: C

Explanation:

You can use the AWS Policy Generator to create a bucket policy for your Amazon S3 bucket. You can then use the generated document to set your bucket policy by using the Amazon S3 console, by a number of third-party tools, or via your application.

You use an Amazon S3 bucket policy that specifies a wildcard (*) in the principal element, which means anyone can access the bucket. With anonymous access, anyone (including users without an AWS account) will be able to access the bucket.

Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/iam-troubleshooting.html#d0e20565>

NEW QUESTION 190

What is the time period with which metric data is sent to CloudWatch when detailed monitoring is enabled on an Amazon EC2 instance?

- A. 15 minutes
- B. 5 minutes
- C. 1 minute
- D. 45 seconds

Answer: C

Explanation:

By default, Amazon EC2 metric data is automatically sent to CloudWatch in 5-minute periods. However, you can, enable detailed monitoring on an Amazon EC2 instance, which sends data to CloudWatch in

1-minute periods

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

NEW QUESTION 194

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

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

- C. Data persists.
- D. Data is delete

Answer: D

Explanation:

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

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

NEW QUESTION 197

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

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

Answer: A

Explanation:

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

Amazon CloudFront billing is mainly affected by Data Transfer Out

Edge Location Traffic Distribution Requests

Dedicated IP SSL Certificates

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

NEW QUESTION 202

Your company has multiple IT departments, each with their own VPC. Some VPCs are located within the same AWS account, and others in a different AWS account. You want to peer together all VPCs to enable the IT departments to have full access to each others' resources. There are certain limitations placed on VPC peering. Which of the following statements is incorrect in relation to VPC peering?

- A. Private DNS values cannot be resolved between instances in peered VPCs.
- B. You can have up to 3 VPC peering connections between the same two VPCs at the same time.
- C. You cannot create a VPC peering connection between VPCs in different regions.
- D. You have a limit on the number active and pending VPC peering connections that you can have per VPC.

Answer: B

Explanation:

To create a VPC peering connection with another VPC, you need to be aware of the following limitations and rules:

You cannot create a VPC peering connection between VPCs that have matching or overlapping CIDR blocks.

You cannot create a VPC peering connection between VPCs in different regions.

You have a limit on the number active and pending VPC peering connections that you can have per VPC. VPC peering does not support transitive peering relationships; in a VPC peering connection, your VPC will not have access to any other VPCs that the peer VPC may be peered with. This includes VPC peering connections that are established entirely within your own AWS account.

You cannot have more than one VPC peering connection between the same two VPCs at the same time. The Maximum Transmission Unit (MTU) across a VPC peering connection is 1500 bytes.

A placement group can span peered VPCs; however, you will not get full-bisection bandwidth between instances in peered VPCs.

Unicast reverse path forwarding in VPC peering connections is not supported.

You cannot reference a security group from the peer VPC as a source or destination for ingress or egress rules in your security group. Instead, reference CIDR blocks of the peer VPC as the source or destination of your security group's ingress or egress rules.

Private DNS values cannot be resolved between instances in peered VPCs. Reference:

<http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/vpc-peering-overview.html#vpc-peering-limitations>

NEW QUESTION 207

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 210

A user has configured a website and launched it using the Apache web server on port 80. The user is using ELB with the EC2 instances for Load Balancing. What should the user do to ensure that the EC2 instances accept requests only from ELB?

- A. Configure the security group of EC2, which allows access to the ELB source security group

- B. Configure the EC2 instance so that it only listens on the ELB port
- C. Open the port for an ELB static IP in the EC2 security group
- D. Configure the security group of EC2, which allows access only to the ELB listener

Answer: A

Explanation:

When a user is configuring ELB and registering the EC2 instances with it, ELB will create a source security group. If the user wants to allow traffic only from ELB, he should remove all the rules set for the other requests and open the port only for the ELB source security group.

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/using-elb-security-groups.html>

NEW QUESTION 212

A user is planning a highly available application deployment with EC2. Which of the below mentioned options will not help to achieve HA?

- A. Elastic IP address
- B. PIOPS
- C. AMI
- D. Availability Zones

Answer: B

Explanation:

In Amazon Web Service, the user can achieve HA by deploying instances in multiple zones. The elastic IP helps the user achieve HA when one of the instances is down but still keeps the same URL. The AM helps launching the new instance. The PIOPS is for the performance of EBS and does not help for HA. Reference:

http://media.amazonwebservices.com/AWS_Web_Hosting_Best_Practices.pdf

NEW QUESTION 214

You need to create a management network using network interfaces for a virtual private cloud (VPC) network. Which of the following statements is incorrect pertaining to Best Practices for Configuring Network Interfaces.

- A. You can detach secondary (ethN) network interfaces when the instance is running or stoppe
- B. However, you can't detach the primary (eth0) interface.
- C. Launching an instance with multiple network interfaces automatically configures interfaces, private IP addresses, and route tables on the operating system of the instance.
- D. You can attach a network interface in one subnet to an instance in another subnet in the same VPC, however, both the network interface and the instance must reside in the same Availability Zone.
- E. Attaching another network interface to an instance is a valid method to increase or double the network bandwidth to or from the dual-homed instance

Answer: D

Explanation:

Best Practices for Configuring Network Interfaces

You can attach a network interface to an instance when it's running (hot attach), when it's stopped (warm attach), or when the instance is being launched (cold attach).

You can detach secondary (ethN) network interfaces when the instance is running or stopped. However, you can't detach the primary (eth0) interface.

You can attach a network interface in one subnet to an instance in another subnet in the same VPC, however, both the network interface and the instance must reside in the same Availability Zone.

When launching an instance from the CLI or API, you can specify the network interfaces to attach to the instance for both the primary (eth0) and additional network interfaces.

Launching an instance with multiple network interfaces automatically configures interfaces, private IP addresses, and route tables on the operating system of the instance.

A warm or hot attach of an additional network interface may require you to manually bring up the second interface, configure the private IP address, and modify the route table accordingly. (Instances running Amazon Linux automatically recognize the warm or hot attach and configure themselves.)

Attaching another network interface to an instance is not a method to increase or double the network bandwidth to or from the dual-homed instance.

Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html#use-network-and-security-applia nces-in-your-vpc>

NEW QUESTION 216

All Amazon EC2 instances are assigned two IP addresses at launch. Which are those?

- A. 2 Elastic IP addresses
- B. A private IP address and an Elastic IP address
- C. A public IP address and an Elastic IP address
- D. A private IP address and a public IP address

Answer: D

Explanation:

In Amazon EC2-Classical every instance is given two IP Addresses: a private IP address and a public IP address

Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-instance-addressing.html#differences>

NEW QUESTION 221

You want to establish a dedicated network connection from your premises to AWS in order to save money by transferring data directly to AWS rather than through your internet service provider. You are sure there must be some other benefits beyond cost savings. Which of the following would not be considered a benefit if you were to establish such a connection?

- A. Elasticity

- B. Compatibility with all AWS services.
- C. Private connectMty to your Amazon VPC.
- D. Everything listed is a benefi

Answer: D

Explanation:

AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS.

Using AWS Direct Connect, you can establish private connectMty between AWS and your datacenter, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections.

You could expect the following benefits if you use AWS Direct Connect. Reduced bandwidth costs

Consistent network performance Compatibility with all AWS services Private connectMty to your Amazon VPC Elasticity

Simplicity

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

NEW QUESTION 224

You are in the process of moving your friend's WordPress site onto AWS to try and save him some money, and you have told him that he should probably also move his domain name. He asks why he can't leave his domain name where it is and just have his infrastructure on AWS. What would be an incorrect response to his question ?

- A. Route 53 offers low query latency for your end users.
- B. Route 53 is designed to automatically answer queries from the optimal location depending on network conditions.
- C. The globally distributed nature of AWS's DNS servers helps ensure a consistent ability to route your end users to your application.
- D. Route 53 supports Domain Name System Security Extensions (DNSSEC).

Answer: D

Explanation:

Amazon Route 53 provides highly available and scalable Domain Name System (DNS), domain name registration, and health-checking web services.

Route 53 is built using AWS's highly available and reliable infrastructure. The globally distributed nature of our DNS servers helps ensure a consistent ability to route your end users to your application by circumventing any internet or network related issues. Route 53 is designed to provide the level of dependability required by important applications. Using a global anycast network of DNS servers around the world, Route 53 is designed to automatically answer queries from the optimal location depending on network conditions. As a result, the service offers low query latency for your end users.

Amazon Route 53 does not support Domain Name System Security Extensions (DNSSEC) at this time. Reference: <https://aws.amazon.com/route53/faqs/>

NEW QUESTION 226

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

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

Answer: A

Explanation:

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

NEW QUESTION 231

While controlling access to Amazon EC2 resources, which of the following acts as a firewall that controls the traffic allowed to reach one or more instances?

- A. A security group
- B. An instance type
- C. A storage cluster
- D. An object

Answer: A

Explanation:

A security group acts as a firewall that controls the traffic allowed to reach one or more instances. When you launch an instance, you assign it one or more security groups.

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

NEW QUESTION 234

Just when you thought you knew every possible storage option on AWS you hear someone mention Reduced Redundancy Storage (RRS) within Amazon S3. What is the ideal scenario to use Reduced Redundancy Storage (RRS)?

- A. Huge volumes of data
- B. Sensitive data
- C. Non-critical or reproducible data
- D. Critical data

Answer: C

Explanation:

Reduced Redundancy Storage (RRS) is a new storage option within Amazon S3 that enables customers to reduce their costs by storing non-critical, reproducible data at lower levels of redundancy than Amazon S3's standard storage. RRS provides a lower cost, less durable, highly available storage option that is designed

to sustain the loss of data in a single facility.

RRS is ideal for non-critical or reproducible data.

For example, RRS is a cost-effective solution for sharing media content that is durably stored elsewhere. RRS also makes sense if you are storing thumbnails and other resized images that can be easily reproduced from an original image.

Reference: <https://aws.amazon.com/s3/faqs/>

NEW QUESTION 239

Your manager has come to you saying that he is very confused about the bills he is receiving from AWS as he is getting different bills for every user and needs you to look into making it more understandable. Which of the following would be the best solution to meet his request?

- A. AWS Billing Aggregation
- B. Consolidated Billing
- C. Deferred Billing
- D. Aggregated Billing

Answer: B

Explanation:

Consolidated Billing enables you to consolidate payment for multiple AWS accounts within your company by designating a single paying account. Consolidated Billing enables you to see a combined view of AWS costs incurred by all accounts, as well as obtain a detailed cost report for each of the individual AWS accounts associated with your "Paying Account". Consolidated Billing is offered at no additional charge. Reference: <https://aws.amazon.com/billing/faqs/>

NEW QUESTION 244

You have been asked to set up a database in AWS that will require frequent and granular updates. You know that you will require a reasonable amount of storage space but are not sure of the best option. What is the recommended storage option when you run a database on an instance with the above criteria?

- A. Amazon S3
- B. Amazon EBS
- C. AWS Storage Gateway
- D. Amazon Glacier

Answer: B

Explanation:

Amazon EBS provides durable, block-level storage volumes that you can attach to a running Amazon EC2 instance. You can use Amazon EBS as a primary storage device for data that requires frequent and granular updates. For example, Amazon EBS is the recommended storage option when you run a database on an instance.

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

NEW QUESTION 248

A user has hosted an application on EC2 instances. The EC2 instances are configured with ELB and Auto Scaling. The application server session time out is 2 hours. The user wants to configure connection draining to ensure that all in-flight requests are supported by ELB even though the instance is being deregistered. What time out period should the user specify for connection draining?

- A. 1 hour
- B. 30 minutes
- C. 5 minutes
- D. 2 hours

Answer: A

Explanation:

The Elastic Load Balancer connection draining feature causes the load balancer to stop sending new requests to the back-end instances when the instances are deregistering or become unhealthy, while ensuring that in-flight requests continue to be served. The user can specify a maximum time of 3600 seconds (1 hour) for the load balancer to keep the connections alive before reporting the instance as deregistered. If the user does not specify the maximum timeout period, by default, the load balancer will close the connections to the deregistering instance after 300 seconds.

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/config-conn-drain.html>

NEW QUESTION 250

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 255

You have a Business support plan with AWS. One of your EC2 instances is running Microsoft Windows Server 2008 R2 and you are having problems with the

software. Can you receive support from AWS for this software?

- A. Yes
- B. No, AWS does not support any third-party software.
- C. No, Microsoft Windows Server 2008 R2 is not supported.
- D. No, you need to be on the enterprise support plan

Answer: A

Explanation:

Third-party software support is available only to AWS Support customers enrolled for Business or Enterprise Support. Third-party support applies only to software running on Amazon EC2 and does not extend to assisting with on-premises software. An exception to this is a VPN tunnel configuration running supported devices for Amazon VPC.

Reference: <https://aws.amazon.com/premiumsupport/features/>

NEW QUESTION 258

After deciding that EMR will be useful in analysing vast amounts of data for a gaming website that you are architecting you have just deployed an Amazon EMR Cluster and wish to monitor the cluster performance. Which of the following tools cannot be used to monitor the cluster performance?

- A. Kinesis
- B. Ganglia
- C. CloudWatch Metrics
- D. Hadoop Web Interfaces

Answer: A

Explanation:

Amazon EMR provides several tools to monitor the performance of your cluster. Hadoop Web Interfaces

Every cluster publishes a set of web interfaces on the master node that contain information about the cluster. You can access these web pages by using an SSH tunnel to connect them on the master node. For more information, see [View Web Interfaces Hosted on Amazon EMR Clusters](#).

CloudWatch Metrics

Every cluster reports metrics to CloudWatch. CloudWatch is a web service that tracks metrics, and which you can use to set alarms on those metrics. For more information, see [Monitor Metrics with CloudWatch](#). Ganglia

Ganglia is a cluster monitoring tool. To have this available, you have to install Ganglia on the cluster when you launch it. After you've done so, you can monitor the cluster as it runs by using an SSH tunnel to connect to the Ganglia UI running on the master node. For more information, see [Monitor Performance with Ganglia](#).

Reference:

<http://docs.aws.amazon.com/ElasticMapReduce/latest/DeveloperGuide/emr-troubleshoot-tools.html>

NEW QUESTION 260

A user has launched one EC2 instance in the US West region. The user wants to access the RDS instance launched in the US East region from that EC2 instance. How can the user configure the access for that EC2 instance?

- A. Configure the IP range of the US West region instance as the ingress security rule of RDS
- B. It is not possible to access RDS of the US East region from the US West region
- C. Open the security group of the US West region in the RDS security group's ingress rule
- D. Create an IAM role which has access to RDS and launch an instance in the US West region with it

Answer: A

Explanation:

The user cannot authorize an Amazon EC2 security group if it is in a different AWS Region than the RDS DB instance. The user can authorize an IP range or specify an Amazon EC2 security group in the same region that refers to an IP address in another region.

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

NEW QUESTION 262

You need to create a load balancer in a VPC network that you are building. You can make your load balancer internal (private) or internet-facing (public). When you make your load balancer internal, a DNS name will be created, and it will contain the private IP address of the load balancer. An internal load balancer is not exposed to the internet. When you make your load balancer internet-facing, a DNS name will be created with the public IP address. If you want the Internet-facing load balancer to be connected to the Internet, where must this load balancer reside?

- A. The load balancer must reside in a subnet that is connected to the internet using the internet gateway.
- B. The load balancer must reside in a subnet that is not connected to the internet.
- C. The load balancer must not reside in a subnet that is connected to the internet.
- D. The load balancer must be completely outside of your VPC

Answer: A

Explanation:

When you create an internal Elastic Load Balancer in a VPC, you need to select private subnets that are in the same Availability Zone as your instances. If the VPC Elastic Load Balancer is to be public facing, you need to create the Elastic Load Balancer in a public subnet. A subnet is a public subnet if it is attached to an Internet Gateway (IGW) with a defined route to that gateway. Selecting more than one public subnet increases the availability of your Elastic Load Balancer.

NB - Elastic Load Balancers in EC2-Classic are always Internet-facing load balancers. Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-internet-facing-load-balancers.html>

NEW QUESTION 267

You need to set up a security certificate for a client's e-commerce website as it will use the HTTPS protocol. Which of the below AWS services do you need to access to manage your SSL server certificate?

- A. AWS Directory Service

- B. AWS Identity & Access Management
- C. AWS CloudFormation
- D. Amazon Route 53

Answer: B

Explanation:

AWS Identity and Access Management (IAM) is a web service that enables Amazon Web Services (AWS) customers to manage users and user permissions in AWS.

All your SSL server certificates are managed by AWS Identity and Access management (IAM). Reference:
<http://docs.aws.amazon.com/IAM/latest/UserGuide/ManagingServerCerts.html>

NEW QUESTION 271

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 272

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

- A. Web sewers: store read-only data in 53, and copy from 53 to root volume at boot tim
- B. App servers: share state using a combination of DynamoDB and IP unicas
- C. Database: use RDS with multi-AZ deployment and one or more read replica
- D. Backup: web servers, app servers, and database backed up weekly to Glacier using snapshots.
- E. Web sewers: store read-only data in an EC2 NFS server, mount to each web server at boot tim
- F. App servers: share state using a combination of DynamoDB and IP multicas
- G. Database: use RDS with multi-AZ deployment and one or more Read Replica
- H. Backup: web and app servers backed up weekly via AM Is, database backed up via DB snapshots.
- I. Web servers: store read-only data in 53, and copy from 53 to root volume at boot tim
- J. App servers: share state using a combination of DynamoDB and IP unicas
- K. Database: use RDS with multi-AZ deployment and one or more Read Replica
- L. Backup: web and app servers backed up weekly viaAM Is, database backed up via DB snapshots.
- M. Web servers: store read-only data in 53, and copy from 53 to root volume at boot tim
- N. App servers: share state using a combination of DynamoDB and IP unicas
- O. Database: use RDS with multi-AZ deploymen
- P. Backup: web and app sewers backed up weekly via ANI Is, database backed up via DB snapshots.

Answer: C

Explanation:

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

Benefits

Enhanced Durability

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

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

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

Increased Availability

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

The availability benefits of Multi-AZ deployments also extend to planned maintenance and backups.

In the case of system upgrades like OS patching or DB Instance scaling, these operations are applied first on the standby, prior to the automatic failover. As a result, your availability impact is, again, only the time required for automatic fail over to complete.

Unlike Single-AZ deployments, 1/0 actMty is not suspended on your primary during backup for Multi-AZ deployments for the MySQL, Oracle, and PostgreSQL

engines, because the backup is taken from the standby. However, note that you may still experience elevated latencies for a few minutes during backups for Multi-AZ deployments.

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

No Administrative Intervention

DB Instance failover is fully automatic and requires no administrative intervention. Amazon RDS monitors the health of your primary and standbys, and initiates a failover automatically in response to a variety of failure conditions.

Failover conditions

Amazon RDS detects and automatically recovers from the most common failure scenarios for Multi-AZ deployments so that you can resume database operations as quickly as possible without administrative intervention. Amazon RDS automatically performs a failover in the event of any of the following:

Loss of availability in primary Availability Zone Loss of network connectivity to primary Compute unit failure on primary

Storage failure on primary

Note: When operations such as DB Instance scaling or system upgrades like OS patching are initiated for Multi-AZ deployments, for enhanced availability, they are applied first on the standby prior to an automatic failover. As a result, your availability impact is limited only to the time required for automatic failover to complete.

Note that Amazon RDS Multi-AZ deployments do not failover automatically in response to database operations such as long running queries, deadlocks or database corruption errors.

NEW QUESTION 276

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

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

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

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

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

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

Answer: C

NEW QUESTION 277

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

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

Answer: AC

Explanation:

Reference:

http://tech.com/wp-content/themes/optimize/download/AWSDisaster_Recovery.pdf (page 6)

NEW QUESTION 282

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

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

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

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

Answer: A

Explanation:

ElastiCache for Memcached

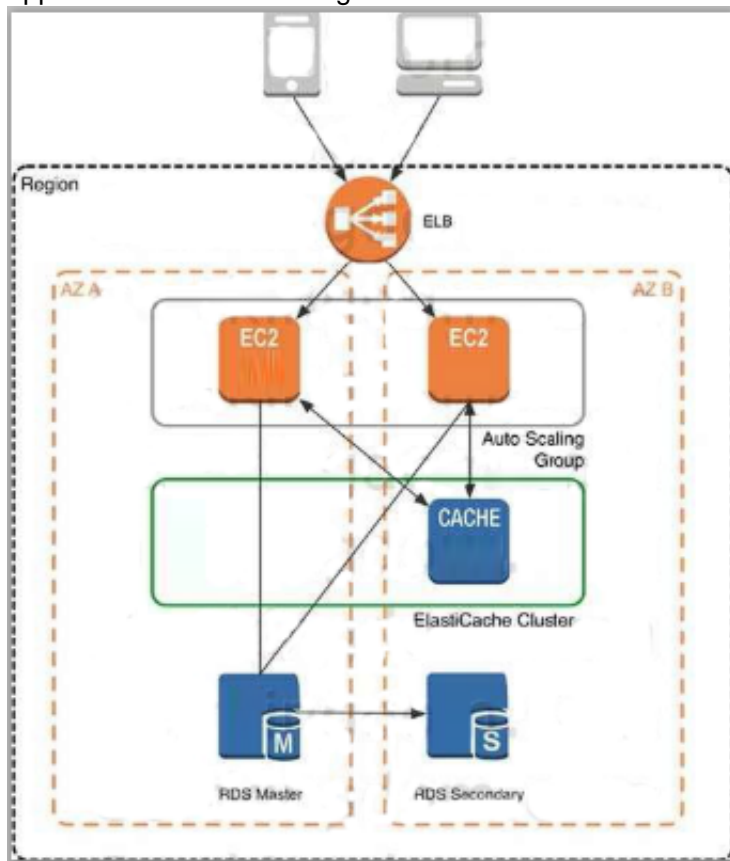
The primary goal of caching is typically to offload reads from your database or other primary data source. In most apps, you have hot spots of data that are regularly queried, but only updated periodically. Think of the front page of a blog or news site, or the top 100 leaderboard in an online game. In this type of case, your app can receive dozens, hundreds, or even thousands of requests for the same data before it's updated again. Having your caching layer handle these queries has several advantages. First, it's considerably cheaper to add an in-memory cache than to scale up to a larger database cluster. Second, an in-memory cache is also easier to scale out, because it's easier to distribute an in-memory cache horizontally than a relational database.

Last, a caching layer provides a request buffer in the event of a sudden spike in usage. If your app or game ends up on the front page of Reddit or the App Store, it's not unheard of to see a spike that is 10 to 100 times your normal application load. Even if you autoscale your application instances, a I/O request spike will likely make your database very unhappy.

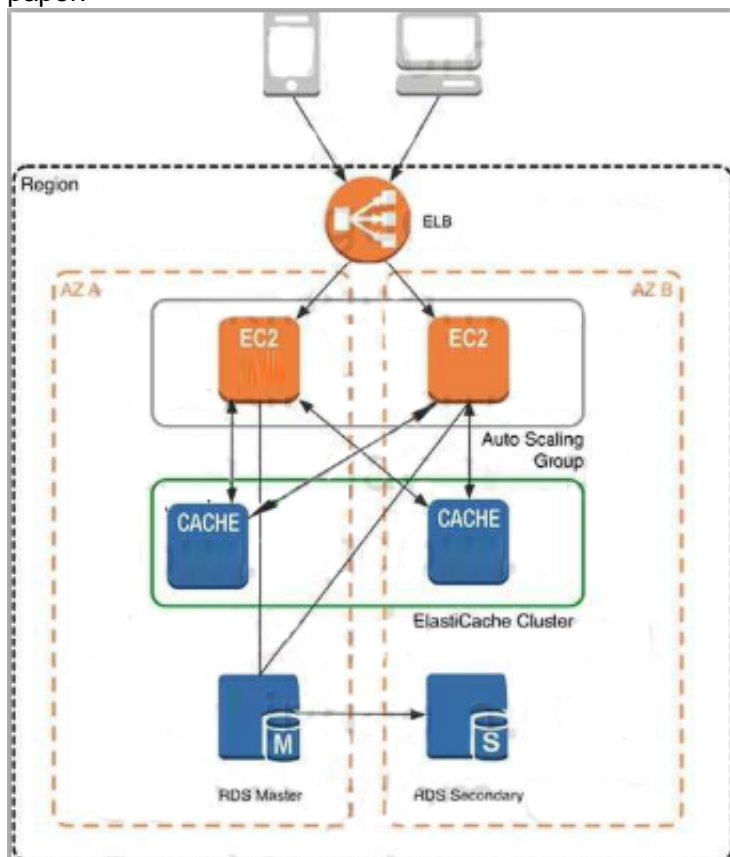
Let's focus on ElastiCache for Memcached first, because it is the best fit for a caching focused solution. We'll revisit Redis later in the paper, and weigh its advantages and disadvantages.

Architecture with ElastiCache for Memcached

When you deploy an ElastiCache Memcached cluster, it sits in your application as a separate tier alongside your database. As mentioned previously, Amazon ElastiCache does not directly communicate with your database tier, or indeed have any particular knowledge of your database. A simplified deployment for a web application looks something like this:



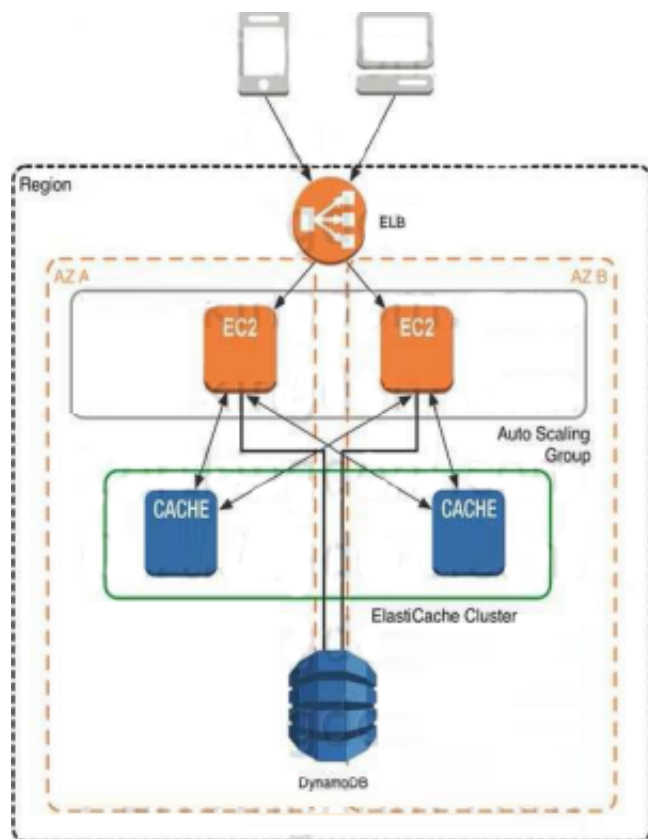
In this architecture diagram, the Amazon EC2 application instances are in an Auto Scaling group, located behind a load balancer using Elastic Load Balancing, which distributes requests among the instances. As requests come into a given EC2 instance, that EC2 instance is responsible for communicating with ElastiCache and the database tier. For development purposes, you can begin with a single ElastiCache node to test your application, and then scale to additional cluster nodes by modifying the ElastiCache cluster. As you add additional cache nodes, the EC2 application instances are able to distribute cache keys across multiple ElastiCache nodes. The most common practice is to use client-side sharding to distribute keys across cache nodes, which we will discuss later in this paper.



When you launch an ElastiCache cluster, you can choose the Availability Zone(s) that the cluster lives in. For best performance, you should configure your cluster to use the same Availability Zones as your application servers. To launch an ElastiCache cluster in a specific Availability Zone, make sure to specify the Preferred Zone(s) option during cache cluster creation. The Availability Zones that you specify will be where ElastiCache will launch your cache nodes. We recommend that you select Spread Nodes Across Zones, which tells ElastiCache to distribute cache nodes across these zones as evenly as possible. This distribution will mitigate the impact of an Availability Zone disruption on your ElastiCache nodes. The trade-off is that some of the requests from your application to ElastiCache will go to a node in a different Availability Zone, meaning latency will be slightly higher.

For more details, refer to Creating a Cache Cluster in the Amazon ElastiCache User Guide.

As mentioned at the outset, ElastiCache can be coupled with a wide variety of databases. Here is an example architecture that uses Amazon DynamoDB instead of Amazon RDS and IVySQL:



This combination of DynamoDB and ElastiCache is very popular with mobile and game companies, because DynamoDB allows for higher write throughput at lower cost than traditional relational databases. In addition, DynamoDB uses a key-value access pattern similar to ElastiCache, which also simplifies the programming model. Instead of using relational SQL for the primary database but then key-value patterns for the cache, both the primary database and cache can be programmed similarly.

In this architecture pattern, DynamoDB remains the source of truth for data, but application reads are offloaded to ElastiCache for a speed boost.

NEW QUESTION 285

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

AWS-Solution-Architect-Associate Practice Exam Features:

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

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