

Amazon

Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate



NEW QUESTION 1

In the context of AWS support, why must an EC2 instance be unreachable for 20 minutes rather than allowing customers to open tickets immediately?

- A. Because most reachability issues are resolved by automated processes in less than 20 minutes
- B. Because all EC2 instances are unreachable for 20 minutes every day when AWS does routine maintenance
- C. Because all EC2 instances are unreachable for 20 minutes when first launched
- D. Because of all the reasons listed here

Answer: A

Explanation:

An EC2 instance must be unreachable for 20 minutes before opening a ticket, because most reachability issues are resolved by automated processes in less than 20 minutes and will not require any action on the part of the customer. If the instance is still unreachable after this time frame has passed, then you should open a case with support.

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

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

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

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

Answer: A

Explanation:

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

Gateway-cached volumes — You store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally. Gateway-cached volumes offer a substantial cost savings on primary storage and minimize the need to scale your storage on-premises. You also retain low-latency access to your frequently accessed data.

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

Reference: <http://docs.aws.amazon.com/storagegateway/latest/userguide/volume-gateway.html>

NEW QUESTION 4

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 5

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 6

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 7

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

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

Answer: A

Explanation:

Amazon RDS gives you access to the capabilities of a familiar MySQL, Oracle, SQL Server, or PostgreSQL database engine. This means that the code, applications, and tools you already use today with your existing databases can be used with Amazon RDS. Amazon RDS automatically patches the database software and backs up your database, storing the backups for a user-defined retention period and enabling point-in-time recovery.
Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html>

NEW QUESTION 8

An edge location refers to which Amazon Web Service?

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

Answer: C

Explanation:

Amazon CloudFront is a content distribution network. A content delivery network or content distribution network (CDN) is a large distributed system of servers deployed in multiple data centers across the world. The location of the data center used for CDN is called edge location.

Amazon CloudFront can cache static content at each edge location. This means that your popular static content (e.g., your site's logo, navigational images, cascading style sheets, JavaScript code, etc.) will be available at a nearby edge location for the browsers to download with low latency and improved performance for viewers. Caching popular static content with Amazon CloudFront also helps you offload requests for such files from your origin 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 9

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 10

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

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

Answer: B

Explanation:

Key pairs consist of a public and private key, where you use the private key to create a digital signature, and then AWS uses the corresponding public key to validate the signature. Key pairs are used only for Amazon EC2 and Amazon CloudFront. Reference: <http://docs.aws.amazon.com/general/latest/gr/aws-sec-cred-types.html>

NEW QUESTION 10

Does Amazon DynamoDB support both increment and decrement atomic operations?

- A. Only increment, since decrement are inherently impossible with DynamoDB's data model.
- B. No, neither increment nor decrement operations.
- C. Yes, both increment and decrement operations.
- D. Only decrement, since increment are inherently impossible with DynamoDB's data mode

Answer: C

Explanation:

Amazon DynamoDB supports increment and decrement atomic operations. Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/APISummary.html>

NEW QUESTION 13

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 acces

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 17

Which of the following is true of Amazon EC2 security group?

- A. You can modify the outbound rules for EC2-Classic.
- B. You can modify the rules for a security group only if the security group controls the traffic for just one instance.
- C. You can modify the rules for a security group only when a new instance is created.
- D. You can modify the rules for a security group at any tim

Answer: D

Explanation:

A security group acts as a virtual firewall that controls the traffic for one or more instances. When you launch an instance, you associate one or more security groups with the instance. You add rules to each security group that allow traffic to or from its associated instances. You can modify the rules for a security group at any time; the new rules are automatically applied to all instances that are associated with the security group.

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

NEW QUESTION 19

In EC2, what happens to the data in an instance store if an instance reboots (either intentionally or unintentionally)?

- A. Data is deleted from the instance store for security reasons.
- B. Data persists in the instance store.
- C. Data is partially present in the instance store.
- D. Data in the instance store will be lost

Answer: B

Explanation:

The data in an instance store persists only during the lifetime of its associated instance. If an instance reboots (intentionally or unintentionally), data in the instance store persists. However, data on instance store volumes is lost under the following circumstances.

Failure of an underlying drive

Stopping an Amazon EBS-backed instance Terminating an instance

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

NEW QUESTION 21

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 24

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

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

Answer: C

Explanation:

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

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

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

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

NEW QUESTION 25

Which of the below mentioned options is not available when an instance is launched by Auto Scaling with EC2 Classic?

- A. Public IP
- B. Elastic IP
- C. Private DNS
- D. Private IP

Answer: B

Explanation:

Auto Scaling supports both EC2 classic and EC2-VPC. When an instance is launched as a part of EC2 classic, it will have the public IP and DNS as well as the private IP and DNS.

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

NEW QUESTION 26

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

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

Answer: B

Explanation:

Read Replicas make it easy to take advantage of MySQL's built-in replication functionality to elastically scale out beyond the capacity constraints of a single DB Instance for read-heavy database workloads. There are a variety of scenarios where deploying one or more Read Replicas for a given source DB Instance may make sense. Common reasons for deploying a Read Replica include:

Scaling beyond the compute or I/O capacity of a single DB Instance for read-heavy database workloads. This excess read traffic can be directed to one or more Read Replicas.

Serving read traffic while the source DB Instance is unavailable. If your source DB Instance cannot take I/O requests (e.g. due to I/O suspension for backups or scheduled maintenance), you can direct read traffic to your Read Replica(s). For this use case, keep in mind that the data on the Read Replica may be "stale" since the source DB Instance is unavailable.

Business reporting or data warehousing scenarios; you may want business reporting queries to run against a Read Replica, rather than your primary, production DB Instance.

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

NEW QUESTION 28

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 30

In Amazon EC2 Container Service, are other container types supported?

- A. Yes, EC2 Container Service supports any container service you need.
- B. Yes, EC2 Container Service also supports Microsoft container service.
- C. No, Docker is the only container platform supported by EC2 Container Service presently.
- D. Yes, EC2 Container Service supports Microsoft container service and Openstac

Answer: C

Explanation:

In Amazon EC2 Container Service, Docker is the only container platform supported by EC2 Container Service presently.

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

NEW QUESTION 35

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 39

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

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

Answer: D

Explanation:

The resource record sets contained in a hosted zone must share the same suffix. For example, the example.com hosted zone can contain resource record sets for www.example.com and www.aws.example.com subdomains, but it cannot contain resource record sets for a www.example.ca subdomain.

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

NEW QUESTION 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 45

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 46

You are architecting an auto-scalable batch processing system using video processing pipelines and Amazon Simple Queue Service (Amazon SQS) for a customer. You are unsure of the limitations of SQS and need to find out. What do you think is a correct statement about the limitations of Amazon SQS?

- A. It supports an unlimited number of queues but a limited number of messages per queue for each user but automatically deletes messages that have been in the queue for more than 4 weeks.
- B. It supports an unlimited number of queues and unlimited number of messages per queue for each user but automatically deletes messages that have been in the queue for more than 4 days.
- C. It supports an unlimited number of queues but a limited number of messages per queue for each user but automatically deletes messages that have been in the queue for more than 4 days.
- D. It supports an unlimited number of queues and unlimited number of messages per queue for each user but automatically deletes messages that have been in the queue for more than 4 weeks.

Answer: B

Explanation:

Amazon Simple Queue Service (Amazon SQS) is a messaging queue service that handles message or workflows between other components in a system. Amazon SQS supports an unlimited number of queues and unlimited number of messages per queue for each user. Please be aware that Amazon SQS automatically deletes messages that have been in the queue for more than 4 days.

Reference: <http://aws.amazon.com/documentation/sqs/>

NEW QUESTION 47

An online gaming site asked you if you can deploy a database that is a fast, highly scalable NoSQL database service in AWS for a new site that he wants to build. Which database should you recommend?

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

Answer: A

Explanation:

Amazon DynamoDB is ideal for database applications that require very low latency and predictable performance at any scale but don't need complex querying capabilities like joins or transactions. Amazon DynamoDB is a fully-managed NoSQL database service that offers high performance, predictable throughput and low cost. It is easy to set up, operate, and scale.

With Amazon DynamoDB, you can start small, specify the throughput and storage you need, and easily scale your capacity requirements on the fly. Amazon DynamoDB automatically partitions data over a number of servers to meet your request capacity. In addition, DynamoDB automatically replicates your data synchronously across multiple Availability Zones within an AWS Region to ensure high-availability and data durability.

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

NEW QUESTION 51

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 56

You log in to IAM on your AWS console and notice the following message. "Delete your root access keys." Why do you think IAM is requesting this?

- A. Because the root access keys will expire as soon as you log out.
- B. Because the root access keys expire after 1 week.
- C. Because the root access keys are the same for all users.
- D. Because they provide unrestricted access to your AWS resource

Answer: D

Explanation:

In AWS an access key is required in order to sign requests that you make using the command-line interface (CLI), using the AWS SDKs, or using direct API calls. Anyone who has the access key for your root account has unrestricted access to all the resources in your account, including billing information. One of the best ways to protect your account is to not have an access key for your root account. We recommend that unless you must have a root access key (this is very rare), that you do not generate one. Instead, AWS best practice is to create one or more AWS Identity and Access Management (IAM) users, give them the necessary permissions, and use IAM users for everyday interaction with AWS.

Reference:

<http://docs.aws.amazon.com/general/latest/gr/aws-access-keys-best-practices.html#root-password>

NEW QUESTION 60

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 ENIR 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 65

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.html#CHAP_Troubleshooting.Security

NEW QUESTION 69

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 73

A user has created photo editing software and hosted it on EC2. The software accepts requests from the user about the photo format and resolution and sends a message to S3 to enhance the picture accordingly. Which of the below mentioned AWS services will help make a scalable software with the AWS infrastructure in this scenario?

- A. AWS Simple Notification Service
- B. AWS Simple Queue Service
- C. AWS Elastic Transcoder
- D. AWS Glacier

Answer: B

Explanation:

Amazon Simple Queue Service (SQS) is a fast, reliable, scalable, and fully managed message queuing service. SQS provides a simple and cost-effective way to decouple the components of an application. The user can configure SQS, which will decouple the call between the EC2 application and S3. Thus, the application does not keep waiting for S3 to provide the data.

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

NEW QUESTION 78

Which one of the following answers is not a possible state of Amazon CloudWatch Alarm?

- A. INSUFFICIENT_DATA
- B. ALARM
- C. OK
- D. STATUS_CHECK_FAILED

Answer: D

Explanation:

Amazon CloudWatch Alarms have three possible states: OK: The metric is within the defined threshold ALARM: The metric is outside of the defined threshold

INSUFFICIENT_DATA: The alarm has just started, the metric is not available, or not enough data is available for the metric to determine the alarm state

Reference: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/AlarmThatSendsEmail.html>

NEW QUESTION 79

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 84

A user is planning to launch a scalable web application. Which of the below mentioned options will not affect the latency of the application?

- A. Region.
- B. Provisioned IOPS.
- C. Availability Zone.
- D. Instance size.

Answer: C

Explanation:

In AWS, the instance size decides the I/O characteristics. The provisioned IOPS ensures higher throughput, and lower latency. The region does affect the latency; latency will always be less when the instance is near to the end user. Within a region the user uses any AZ and this does not affect the latency. The AZ is mainly for fault toleration or HA.

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

NEW QUESTION 86

An organization has a statutory requirement to protect the data at rest for data stored in EBS volumes. Which of the below mentioned options can the organization use to achieve data protection?

- A. Data replication.
- B. Data encryption.
- C. Data snapshot.
- D. All the options listed here.

Answer: D

Explanation:

For protecting the Amazon EBS data at REST, the user can use options, such as Data Encryption (Windows / Linux / third party based), Data Replication (AWS internally replicates data for redundancy), and Data Snapshot (for point in time backup).

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

NEW QUESTION 90

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 95

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, you 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 98

You have a number of image files to encode. In an Amazon SQS worker queue, you create an Amazon SQS message for each file specifying the command (jpeg-encode) and the location of the file in Amazon S3. Which of the following statements best describes the functionality of Amazon SQS?

- A. Amazon SQS is a distributed queuing system that is optimized for horizontal scalability, not for single-threaded sending or receiving speeds.
- B. Amazon SQS is for single-threaded sending or receiving speeds.
- C. Amazon SQS is a non-distributed queuing system.
- D. Amazon SQS is a distributed queuing system that is optimized for vertical scalability and for single-threaded sending or receiving speeds.

Answer: A

Explanation:

Amazon SQS is a distributed queuing system that is optimized for horizontal scalability, not for single-threaded sending or receiving speeds. A single client can send or receive Amazon SQS messages at a rate of about 5 to 50 messages per second. Higher receive performance can be achieved by requesting multiple messages (up to 10) in a single call. It may take several seconds before a message that has been to a queue is available to be received.

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

NEW QUESTION 101

A user has set up the CloudWatch alarm on the CPU utilization metric at 50%, with a time interval of 5 minutes and 10 periods to monitor. What will be the state of the alarm at the end of 90 minutes, if the CPU utilization is constant at 80%?

- A. ALERT
- B. ALARM
- C. OK
- D. INSUFFICIENT_DATA

Answer: B

Explanation:

In this case the alarm watches a metric every 5 minutes for 10 intervals. Thus, it needs at least 50 minutes to come to the "OK" state. Till then it will be in the INSUFFICIENT_DATA state.

Since 90 minutes have passed and CPU utilization is at 80% constant, the state of alarm will be "ALARM". Reference: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/AlarmThatSendsEmail.html>

NEW QUESTION 105

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 109

A user is planning to make a mobile game which can be played online or offline and will be hosted on EC2.

The user wants to ensure that if someone breaks the highest score or they achieve some milestone they can inform all their colleagues through email. Which of the below mentioned AWS services helps achieve this goal?

- A. AWS Simple Workflow Service.
- B. AWS Simple Email Service.
- C. Amazon Cognito
- D. AWS Simple Queue Service

Answer: B

Explanation:

Amazon Simple Email Service (Amazon SES) is a highly scalable and cost-effective email-sending service for businesses and developers. It integrates with other AWS services, making it easy to send emails from applications that are hosted on AWS.

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

NEW QUESTION 113

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 114

A user has deployed an application on his private cloud. The user is using his own monitoring tool. He wants to configure it so that whenever there is an error, the monitoring tool will notify him via SMS. Which of the below mentioned AWS services will help in this scenario?

- A. AWS SES
- B. AWS SNS
- C. None because the user infrastructure is in the private cloud.
- D. AWS SMS

Answer: B

Explanation:

Amazon Simple Notification Service (Amazon SNS) is a fast, flexible, and fully managed push messaging service. Amazon SNS can be used to make push notifications to mobile devices. Amazon SNS can deliver notifications by SMS text message or email to the Amazon Simple Queue Service (SQS) queues or to any HTTP endpoint. In this case user can use the SNS apis to send SMS.

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

NEW QUESTION 115

You are configuring a new VPC for one of your clients for a cloud migration project, and only a public VPN will be in place. After you created your VPC, you created a new subnet, a new internet gateway, and attached your internet gateway to your VPC. When you launched your first instance into your VPC, you realized that you aren't able to connect to the instance, even if it is configured with an elastic IP. What should be done to access the instance?

- A. A route should be created as 0.0.0.0/0 and your internet gateway as target.
- B. Attach another ENI to the instance and connect via new ENI.
- C. A NAT instance should be created and all traffic should be forwarded to NAT instance.
- D. A NACL should be created that allows all outbound traffic

Answer: A

Explanation:

All traffic should be routed via Internet Gateway. So, a route should be created with 0.0.0.0/0 as a source, and your Internet Gateway as your target.

Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario1.html

NEW QUESTION 118

What is the default maximum number of Access Keys per user?

- A. 10
- B. 15
- C. 2
- D. 20

Answer: C

Explanation:

The default maximum number of Access Keys per user is 2.

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

NEW QUESTION 119

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

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

Answer: B

Explanation:

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

NEW QUESTION 123

Which of the following statements is true of tagging an Amazon EC2 resource?

- A. You don't need to specify the resource identifier while terminating a resource.
- B. You can terminate, stop, or delete a resource based solely on its tags.
- C. You can't terminate, stop, or delete a resource based solely on its tags.
- D. You don't need to specify the resource identifier while stopping a resource.

Answer: C

Explanation:

You can assign tags only to resources that already exist. You can't terminate, stop, or delete a resource based solely on its tags; you must specify the resource identifier.

Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Using_Tags.html

NEW QUESTION 124

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 expire.
- 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 expire.

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

NEW QUESTION 129

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 131

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

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

Answer: B

Explanation:

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

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

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

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

NEW QUESTION 134

After setting up some EC2 instances you now need to set up a monitoring solution to keep track of these instances and to send you an email when the CPU hits a certain threshold. Which statement below best describes what thresholds you can set to trigger a CloudWatch Alarm?

- A. Set a target value and choose whether the alarm will trigger when the value is greater than (>), greater than or equal to (>=), less than (<), or less than or equal to (<=) that value.
- B. Thresholds need to be set in IAM not CloudWatch

- C. Only default thresholds can be set you can't choose your own thresholds.
- D. Set a target value and choose whether the alarm will trigger when the value hits this threshold

Answer: A

Explanation:

Amazon CloudWatch is a monitoring service for AWS cloud resources and the applications you run on AWS. You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, and set alarms.

When you create an alarm, you first choose the Amazon CloudWatch metric you want it to monitor. Next, you choose the evaluation period (e.g., five minutes or one hour) and a statistical value to measure (e.g., Average or Maximum).

To set a threshold, set a target value and choose whether the alarm will trigger when the value is greater than (>), greater than or equal to (>=), less than (<), or less than or equal to (<=) that value.

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

NEW QUESTION 139

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 (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 140

When does the billing of an Amazon EC2 system begin?

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

Answer: D

Explanation:

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

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

NEW QUESTION 145

What is the data model of DynamoDB?

- A. Since DynamoDB is schema-less, there is no data model.
- B. "Items", with Keys and one or more Attribute; and "Attribute", with Name and Value.
- C. "Table", a collection of Items; "Items", with Keys and one or more Attribute; and "Attribute", with Name and Value.
- D. "Database", which is a set of "Tables", which is a set of "Items", which is a set of "Attributes".

Answer: C

Explanation:

The data model of DynamoDB is: "Table", a collection of Items;

"Items", with Keys and one or more Attribute; "Attribute", with Name and Value.

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

NEW QUESTION 150

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 152

A favored client needs you to quickly deploy a database that is a relational database service with minimal administration as he wants to spend the least amount of time administering it. Which database would be the best option?

- A. Amazon SimpleDB
- B. Your choice of relational AMs on Amazon EC2 and EBS.
- C. Amazon RDS
- D. Amazon Redshift

Answer: C

Explanation:

Amazon Relational Database Service (Amazon RDS) is a web 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.

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

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

NEW QUESTION 154

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 159

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

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

Answer: C

Explanation:

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

NEW QUESTION 161

A for a VPC is a collection of subnets (typically private) that you may want to designate for your backend RDS DB Instances.

- A. DB Subnet Set
- B. RDS Subnet Group
- C. DB Subnet Group
- D. DB Subnet Collection

Answer: C

Explanation:

DB Subnet Groups are a set of subnets (one per Availability Zone of a particular region) designed for your DB instances that reside in a VPC. They make easy to

manage Multi-AZ deployments as well as the conversion from a Single-AZ to a Multi-AZ one.
Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.RDSVPC.html>

NEW QUESTION 163

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 wildcard (*) 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 168

In Amazon CloudFront, if you use Amazon EC2 instances and other custom origins with CloudFront, it is recommended to .

- A. not use Elastic Load Balancing
- B. restrict Internet communication to private instances while allowing outgoing traffic
- C. enable access key rotation for CloudWatch metrics
- D. specify the URL of the load balancer for the domain name of your origin server

Answer: D

Explanation:

In Amazon CloudFront, you should use an Elastic Load Balancing load balancer to handle traffic across multiple Amazon EC2 instances and to isolate your application from changes to Amazon EC2 instances. When you create your CloudFront distribution, specify the URL of the load balancer for the domain name of your origin server.

Reference: <http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/CustomOriginBestPractices.html>

NEW QUESTION 170

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 174

A user is sending bulk emails using AWS SES. The emails are not reaching some of the targeted audience because they are not authorized by the ISPs. How can the user ensure that the emails are all delivered?

- A. Send an email using DKIM with SES.
- B. Send an email using SMTP with SES.
- C. Open a ticket with AWS support to get it authorized with the ISP.
- D. Authorize the ISP by sending emails from the development account

Answer: A

Explanation:

Domain Keys Identified Mail (DKIM) is a standard that allows senders to sign their email messages and ISPs, and use those signatures to verify that those messages are legitimate and have not been modified by a third party in transit.

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

NEW QUESTION 176

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 178

A user is trying to launch a similar EC2 instance from an existing instance with the option "Launch More like this". The AMI of the selected instance is deleted. What will happen in this case?

- A. AWS does not need an AMI for the "Launch more like this" option
- B. AWS will launch the instance but will not create a new AMI
- C. AWS will create a new AMI and launch the instance
- D. AWS will throw an error saying that the AMI is deregistered

Answer: D

Explanation:

If the user has deregistered the AMI of an EC2 instance and is trying to launch a similar instance with the option "Launch more like this", AWS will throw an error saying that the AMI is deregistered or not available.

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

NEW QUESTION 180

A company wants to review the security requirements of Glacier. Which of the below mentioned statements is true with respect to the AWS Glacier data security?

- A. All data stored on Glacier is protected with AES-256 serverside encryption.
- B. All data stored on Glacier is protected with AES-128 serverside encryption.
- C. The user can set the serverside encryption flag to encrypt the data stored on Glacier.
- D. The data stored on Glacier is not encrypted by default

Answer: A

Explanation:

For Amazon Web Services, all the data stored on Amazon Glacier is protected using serverside encryption. AWS generates separate unique encryption keys for each Amazon Glacier archive, and encrypts it using AES-256. The encryption key then encrypts itself using AES-256 with a master key that is stored in a secure location.

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

NEW QUESTION 181

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 184

You are playing around with setting up stacks using JSON templates in CloudFormation to try and understand them a little better. You have set up about 5 or 6 but now start to wonder if you are being charged for these stacks. What is AWS's billing policy regarding stack resources?

- A. You are not charged for the stack resources if they are not taking any traffic.
- B. You are charged for the stack resources for the time they were operating (even if you deleted the stack right away)
- C. You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 60 minutes)
- D. You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 30 minutes)

Answer: B

Explanation:

A stack is a collection of AWS resources that you can manage as a single unit. In other words, you can create, update, or delete a collection of resources by creating, updating, or deleting stacks. All the resources in a stack are defined by the stack's AWS CloudFormation template. A stack, for instance, can include all the resources required to run a web application, such as a web server, a database, and networking rules. If you no longer require that web application, you can simply delete the stack, and all of its related resources are deleted.

You are charged for the stack resources for the time they were operating (even if you deleted the stack right away).

Reference: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacks.html>

NEW QUESTION 185

In Route 53, what does a Hosted Zone refer to?

- A. A hosted zone is a collection of geographical load balancing rules for Route 53.
- B. A hosted zone is a collection of resource record sets hosted by Route 53.
- C. A hosted zone is a selection of specific resource record sets hosted by CloudFront for distribution to Route 53.
- D. A hosted zone is the Edge Location that hosts the Route 53 records for a use

Answer: B

Explanation:

A Hosted Zone refers to a selection of resource record sets hosted by Route 53.

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

NEW QUESTION 187

Which of the following statements is true of Amazon EC2 security groups?

- A. You can change the outbound rules for EC2-Classi
- B. Also, you can add and remove rules to a group at any time.
- C. You can modify an existing rule in a grou
- D. However, you can't add and remove rules to a group.
- E. None of the statements are correct.
- F. You can't change the outbound rules for EC2-Classi
- G. However, you can add and remove rules to a group at any tim

Answer: D

Explanation:

When dealing with security groups, bear in mind that you can freely add and remove rules from a group, but you can't change the outbound rules for EC2-Classic. If you're using the Amazon EC2 console, you can modify existing rules, and you can copy the rules from an existing security group to a new security group.

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

NEW QUESTION 190

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

Inbound SSH traffic.

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

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 191

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 193

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 197

A user is planning to host a mobile game on EC2 which sends notifications to active users on either high score or the addition of new features. The user should get this notification when he is online on his mobile device. Which of the below mentioned AWS services can help achieve this functionality?

- A. AWS Simple Notification Service.
- B. AWS Simple Email Service.
- C. AWS NMobile Communication Service.
- D. AWS Simple Queue Service.

Answer: A

Explanation:

Amazon Simple Notification Service (Amazon SNS) is a fast, flexible, and 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>

NEW QUESTION 199

You have written a CloudFormation template that creates 1 Elastic Load Balancer fronting 2 EC2 Instances. Which section of the template should you edit so that the DNS of the load balancer is returned upon creation of the stack?

- A. Resources
- B. Outputs
- C. Parameters
- D. Mappings

Answer: B

Explanation:

You can use AWS CloudFormation's sample templates or create your own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run your application.

Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/outputs-section-structure.html>

NEW QUESTION 202

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 205

Identify a true statement about the On-Demand instances purchasing option provided by Amazon EC2.

- A. Pay for the instances that you use by the hour, with no long-term commitments or up-front payments.
- B. Make a low, one-time, up-front payment for an instance, reserve it for a one- or three-year term, and pay a significantly lower hourly rate for these instances.
- C. Pay for the instances that you use by the hour, with long-term commitments or up-front payments.
- D. Make a high, one-time, all-front payment for an instance, reserve it for a one- or three-year term, and pay a significantly higher hourly rate for these instance

Answer: A

Explanation:

On-Demand instances allow you to pay for the instances that you use by the hour, with no long-term commitments or up-front payments.

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

NEW QUESTION 206

Which of the following statements is NOT true about using Elastic IP Address (EIP) in EC2-Classic and EC2-VPC platforms?

- A. In the EC2-VPC platform, the Elastic IP Address (EIP) does not remain associated with the instance when you stop it.

- B. In the EC2-Classic platform, stopping the instance disassociates the Elastic IP Address (EIP) from it.
- C. In the EC2-VPC platform, if you have attached a second network interface to an instance, when you disassociate the Elastic IP Address (EIP) from that instance, a new public IP address is not assigned to the instance automatically; you'll have to associate an EIP with it manually.
- D. In the EC2-Classic platform, if you disassociate an Elastic IP Address (EIP) from the instance, the instance is automatically assigned a new public IP address within a few minutes.

Answer: A

Explanation:

In the EC2-Classic platform, when you associate an Elastic IP Address (EIP) with an instance, the instance's current public IP address is released to the EC2-Classic public IP address pool. If you disassociate an EIP from the instance, the instance is automatically assigned a new public IP address within a few minutes. In addition, stopping the instance also disassociates the EIP from it.

But in the EC2-VPC platform, when you associate an EIP with an instance in a default Virtual Private Cloud (VPC), or an instance in which you assigned a public IP to the eth0 network interface during launch, its current public IP address is released to the EC2-VPC public IP address pool. If you disassociate an EIP from the instance, the instance is automatically assigned a new public IP address within a few minutes. However, if you have attached a second network interface to the instance, the instance is not automatically assigned a new public IP address; you'll have to associate an EIP with it manually. The EIP remains associated with the instance when you stop it.

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

NEW QUESTION 210

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 VP

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 211

Can you move a Reserved Instance from one Availability Zone to another?

- A. Yes, but each Reserved Instance is associated with a specific Region that cannot be changed.
- B. Yes, only in US-West-2.
- C. Yes, only in US-East-1.
- D. No

Answer: A

Explanation:

Each Reserved Instance is associated with a specific Region, which is fixed for the lifetime of the reservation and cannot be changed. Each reservation can, however, be used in any of the available AZs within the associated Region.

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

NEW QUESTION 213

An application hosted at the EC2 instance receives an HTTP request from ELB. The same request has an X-Forwarded-For header, which has three IP addresses. Which system's IP will be a part of this header?

- A. Previous Request IP address.
- B. Client IP address.
- C. All of the answers listed here.
- D. Load Balancer IP address

Answer: C

Explanation:

When a user sends a request to ELB over HTTP/HTTPS, the request header log at the instance will only receive the IP of ELB. This is because ELB is the interceptor between the EC2 instance and the client request. To get the client IP, use the header X-Forwarded-For in header. The client IP address in the X-Forwarded-For request header is followed by the IP addresses of each successive proxy that passes along the request. The last IP address is the IP address that connects to the back-end application instance. e.g. if the HTTP request already has a header when it reaches the Load Balancer, the IP address from which the request came is appended at the end of the header followed by the IP address of the Load Balancer. In such cases, the X-Forwarded-For request header takes the following form:

X-Forwarded-For: clientIPAddress, previousRequestIPAddress, LoadBalancerIPAddress. Reference:

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

NEW QUESTION 216

You need to develop and run some new applications on AWS and you know that Elastic Beanstalk and CloudFormation can both help as a deployment mechanism for a broad range of AWS resources. Which of the following statements best describes the differences between Elastic Beanstalk and CloudFormation?

- A. Elastic Beanstalk uses Elastic load balancing and CloudFormation doesn't.
- B. CloudFormation is faster in deploying applications than Elastic Beanstalk.
- C. Elastic Beanstalk is faster in deploying applications than CloudFormation.
- D. CloudFormation is much more powerful than Elastic Beanstalk, because you can actually design and script custom resources

Answer: D

Explanation:

These services are designed to complement each other. AWS Elastic Beanstalk provides an environment to easily develop and run applications in the cloud. It is integrated with developer tools and provides a one-stop experience for you to manage the lifecycle of your applications. AWS CloudFormation is a convenient deployment mechanism for a broad range of AWS resources. It supports the infrastructure needs of many different types of applications such as existing enterprise applications, legacy applications, applications built using a variety of AWS resources and container-based solutions (including those built using AWS Elastic Beanstalk).

AWS CloudFormation introduces two new concepts: The template, a JSON-format, text-based file that describes all the AWS resources you need to deploy to run your application and the stack, the set of AWS resources that are created and managed as a single unit when AWS CloudFormation instantiates a template.

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

NEW QUESTION 218

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 223

Your company has HQ in Tokyo and branch offices all over the world and is using a logistics software with a multi-regional deployment on AWS in Japan, Europe and USA, The logistic software has a 3- tier architecture and currently uses MySQL 5.6 for data persistence. Each region has deployed its own database In the HQ region you run an hourly batch process reading data from every region to compute cross regional reports that are sent by email to all offices this batch process must be completed as fast as possible to quickly optimize logistics how do you build the database architecture in order to meet the requirements'?

- A. For each regional deployment, use RDS MySQL with a master in the region and a read replica in the HQ region
- B. For each regional deployment, use MySQL on EC2 with a master in the region and send hourly EBS snapshots to the HQ region
- C. For each regional deployment, use RDS MySQL with a master in the region and send hourly RDS snapshots to the HQ region
- D. For each regional deployment, use MySQL on EC2 with a master in the region and use 53 to copy data files hourly to the HQ region
- E. Use Direct Connect to connect all regional MySQL deployments to the HQ region and reduce network latency for the batch process

Answer: A

NEW QUESTION 228

A customer has a 10 GB AWS Direct Connect connection to an AWS region where they have a web application hosted on Amazon Elastic Computer Cloud (EC2). The application has dependencies on an on-premises mainframe database that uses a BASE (Basic Available. Sort stale Eventual consistency) rather than an ACID (Atomicity. Consistency isolation. Durability) consistency model.

The application is exhibiting undesirable behavior because the database is not able to handle the volume of writes. How can you reduce the load on your on-premises database resources in the most cost-effective way?

- A. Use an Amazon Elastic Map Reduce (EMR) S3DistCp as a synchronization mechanism between the on-premises database and a Hadoop cluster on AWS.
- B. Modify the application to write to an Amazon SQS queue and develop a worker process to flush the queue to the on-premises database.
- C. Modify the application to use DynamoDB to feed an EMR cluster which uses a map function to write to the on-premises database.
- D. Provision an RDS read-replica database on AWS to handle the writes and synchronize the two databases using Data Pipeline.

Answer: A

Explanation:

Reference: <https://aws.amazon.com/blogs/aws/category/amazon-elastic-map-reduce/>

NEW QUESTION 232

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

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

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

- A. Utilize 53 to collect the inbound sensor data analyze the data from 53 with a daily scheduled Data Pipeline and save the results to a Redshift Cluster.
- B. Utilize Amazon Kinesis to collect the inbound sensor data, analyze the data with Kinesis clients and save the results to a Red shift cluster using EMR.
- C. Utilize SQS to collect the inbound sensor data analyze the data from SQS with Amazon Kinesis and save the results to a Microsoft SQL Server RDS instance.
- D. Utilize EMR to collect the inbound sensor data, analyze the data from EUR with Amazon Kinesis and save me results to Dynamo DB.

Answer: B

NEW QUESTION 236

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 or 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 IOx 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 MySQL:




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 241

Refer to the architecture diagram above of a batch processing solution using Simple Queue Service (SQS) to set up a message queue between EC2 instances which are used as batch processors. Cloud Watch monitors the number of Job requests (queued messages) and an Auto Scaling group adds or deletes batch servers automatically based on parameters set in Cloud Watch alarms. You can use this architecture to implement which of the following features in a cost effective and efficient manner?



- A. Reduce the overall time for executing jobs through parallel processing by allowing a busy EC2 instance that receives a message to pass it to the next instance in a daisy-chain setup.
- B. Implement fault tolerance against EC2 instance failure since messages would remain in SQS and work can continue with recovery of EC2 instances implement fault tolerance against SQS failure by backing up messages to S3.
- C. Implement message passing between EC2 instances within a batch by exchanging messages through SQS.
- D. Coordinate number of EC2 instances with number of job requests automatically thus Improving cost effectiveness.
- E. Handle high priority jobs before lower priority jobs by assigning a priority metadata field to SQS messages.

Answer: D

Explanation:

Reference:

There are cases where a large number of batch jobs may need processing, and where the jobs may need to be re-prioritized.

For example, one such case is one where there are differences between different levels of services for unpaid users versus subscriber users (such as the time until publication) in services enabling, for example, presentation files to be uploaded for publication from a web browser. When the user uploads a presentation file, the conversion processes, for example, for publication are performed as batch processes on the system side, and the file is published after the conversion. Is it then necessary to be able to assign the level of priority to the batch processes for each type of subscriber.

Explanation of the Cloud Solution/Pattern

A queue is used in controlling batch jobs. The queue need only be provided with priority numbers. Job requests are controlled by the queue, and the job requests in the queue are processed by a batch server. In Cloud computing, a highly reliable queue is provided as a service, which you can use to structure a highly reliable batch system with ease. You may prepare multiple queues depending on priority levels, with job requests put into the queues depending on their priority levels, to apply prioritization to batch processes. The performance (number) of batch servers corresponding to a queue must be in accordance with the priority level thereof.

Implementation

In AWS, the queue service is the Simple Queue Service (SQS). Multiple SQS queues may be prepared to prepare queues for individual priority levels (with a priority queue and a secondary queue).

Moreover, you may also use the message Delayed Send function to delay process execution. Use SQS to prepare multiple queues for the individual priority levels. Place those processes to be executed immediately (job requests) in the high priority queue. Prepare numbers of batch servers, for processing the job requests of the queues, depending on the priority levels.

Queues have a message "Delayed Send" function. You can use this to delay the time for starting a process.

Configuration

Benefits

You can increase or decrease the number of servers for processing jobs to change automatically the processing speeds of the priority queues and secondary queues.

You can handle performance and service requirements through merely increasing or decreasing the number of EC2 instances used in job processing.

Even if an EC2 were to fail, the messages (jobs) would remain in the queue service, enabling processing to be continued immediately upon recovery of the EC2 instance, producing a system that is robust to failure.

Cautions

Depending on the balance between the number of EC2 instances for performing the processes and the number of messages that are queued, there may be cases where processing in the secondary queue may be completed first, so you need to monitor the processing speeds in the primary queue and the secondary queue.

NEW QUESTION 246

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

How would you do this while minimizing costs?

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

Answer: A

Explanation:

Overview of Creating Amazon EBS-Backed AMIs

First, launch an instance from an AMI that's similar to the AMI that you'd like to create. You can connect to your instance and customize it. When the instance is configured correctly, ensure data integrity by stopping the instance before you create an AMI, then create the image. When you create an Amazon EBS-backed AMI, we automatically register it for you. Amazon EC2 powers down the instance before creating the AMI to ensure that everything on the instance is stopped and in a consistent state during the creation process. If you're confident that your instance is in a consistent state appropriate for AMI creation, you can tell Amazon EC2 not to power down and reboot the instance. Some file systems, such as XFS, can freeze and unfreeze actMty, making it safe to create the image without rebooting the instance. During the AMI-creation process, Amazon EC2 creates snapshots of your instance's root volume and any other EBS volumes attached to your instance. If any volumes attached to the instance are encrypted, the new AMI only launches successfully on instances that support Amazon EBS encryption. For more information, see Amazon EBS Encryption. Depending on the size of the volumes, it can take several minutes for the AMI-creation process to complete (sometimes up to 24 hours). You may find it more efficient to create snapshots of your volumes prior to creating your AMI. This way, only small, incremental snapshots need to be created when the AMI is created, and the process completes more quickly (the total time for snapshot creation remains the same). For more information, see Creating an Amazon EBS Snapshot. After the process completes, you have a new AMI and snapshot created from the root volume of the instance. When you launch an instance using the new AMI, we create a new EBS volume for its root volume using the snapshot. Both the AMI and the snapshot incur charges to your account until you delete them. For more information, see Deregistering Your AMI. If you add instance-store volumes or EBS volumes to your instance in addition to the root device volume, the block device mapping for the new AMI contains information for these volumes, and the block device mappings for instances that you launch from the new AMI automatically contain information for these volumes. The instance-store volumes specified in the block device mapping for the new instance are new and don't contain any data from the instance store volumes of the instance you used to create the AMI. The data on EBS volumes persists. For more information, see Block Device Mapping.

NEW QUESTION 247

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

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

Answer: AC

NEW QUESTION 248

You've been brought in as solutions architect to assist an enterprise customer with their migration of an e-commerce platform to Amazon Virtual Private Cloud (VPC) The previous architect has already deployed a 3-tier VPC, The configuration is as follows:
VPC: vpc-2f8bc447 IGW: igw-2d8bc445 NACL: ad-208bc448
Subnets and Route Tables: Web sewers: subnet-258bc44d
Application servers: subnet-248bc44c Database sewers: subnet-9189c6f9 Route Tables:
rtb-218bc449 rtb-238bc44b Associations:
subnet-258bc44d : rtb-218bc449 subnet-248bc44c : rtb-238bc44b subnet-9189c6f9 : rtb-238bc44b
You are now ready to begin deploying EC2 instances into the VPC Web servers must have direct access to the internet Application and database sewers cannot have direct access to the internet.
Which configuration below will allow you the ability to remotely administer your application and database servers, as well as allow these sewers to retrieve updates from the Internet?

- A. Create a bastion and NAT instance in subnet-258bc44d, and add a route from rtb- 238bc44b to the NAT instance.
- B. Add a route from rtb-238bc44b to igw-2d8bc445 and add a bastion and NAT instance within subnet-248bc44c.
- C. Create a bastion and NAT instance in subnet-248bc44c, and add a route from rtb- 238bc44b to subneb258bc44d.
- D. Create a bastion and NAT instance in subnet-258bc44d, add a route from rtb-238bc44b to lgw- 2d8bc445, and a new NACL that allows access between subnet-258bc44d and subnet -248bc44c.

Answer: A

NEW QUESTION 249

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

- A. Configure a NAT instance in your VPC Create a default route via the NAT instance and associate it with all subnets Configure a DNS A record that points to the NAT instance public IP address.
- B. Configure a C|oudFront distribution and configure the origin to point to the private IP addresses of your Web sewers Configure a Route53 CNAME record to your Cloud Front distribution.
- C. Place all your web servers behind EL8 Configure a Route53 CNME to point to the ELB DNS name.
- D. Assign EIPs to all web sewer
- E. Configure a Route53 record set with all EIP
- F. With health checks and DNS failover.
- G. Configure ELB with an EIP Place all your Web servers behind ELB Configure a Route53 A record that points to the EIP.

Answer: CD

NEW QUESTION 252

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

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

Answer: ADF

Explanation:

AWS Direct Connect

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 connectivity between AWS 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.

AWS Direct Connect lets you establish a dedicated network connection between your network and one of the AWS Direct Connect locations. Using industry standard 802.1q VLANs, this dedicated connection can be partitioned into multiple virtual interfaces. This allows you to use the same connection to access public resources such as objects stored in Amazon S3 using public IP address space, and private resources such as Amazon EC2 instances running within an Amazon Virtual Private Cloud (VPC) using private IP space, while maintaining network separation between the public and private environments. Virtual interfaces can be reconfigured at any time to meet your changing needs.

What is AWS Direct Connect?

AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 1 gigabit or 10 gigabit Ethernet fiber-optic cable. One end of the cable is connected to your router, the other to an AWS Direct Connect router. With this connection in place, you can create virtual interfaces directly to the AWS cloud (for example, to Amazon Elastic Compute Cloud (Amazon EC2) and Amazon Simple Storage Service (Amazon S3)) and to Amazon Virtual Private Cloud (Amazon VPC), bypassing Internet service providers in your network path. An AWS Direct Connect location provides access to Amazon Web Services in the region it is associated with, as well as access to other US regions. For example, you can provision a single connection to any AWS Direct Connect location in the US and use it to access public AWS services in all US Regions and AWS GovCloud (US).

The following diagram shows how AWS Direct Connect interfaces with your network.

Requirements

To use AWS Direct Connect, your network must meet one of the following conditions:

Your network is colocated with an existing AWS Direct Connect location. For more information on available AWS Direct Connect locations, go to <http://aws.amazon.com/directconnect/>.

You are working with an AWS Direct Connect partner who is a member of the AWS Partner Network (APN). For a list of AWS Direct Connect partners who can help you connect, go to <http://aws.amazon.com/directconnect>

You are working with an independent service provider to connect to AWS Direct Connect. In addition, your network must meet the following conditions:

Connections to AWS Direct Connect require single mode fiber, 1000BASE-LX (1310nm) for 1 gigabit Ethernet, or 10GBASE-LR (1310nm) for 10 gigabit Ethernet. Auto Negotiation for the port must be disabled. You must support 802.1Q VLANs across these connections.

Your network must support Border Gateway Protocol (BGP) and BGP MD5 authentication. Optionally, you may configure Bidirectional Forwarding Detection (BFD).

To connect to Amazon Virtual Private Cloud (Amazon VPC), you must first do the following: Provide a private Autonomous System Number (ASN). Amazon allocates a private IP address in the 169.x.x.x range to you.

Create a virtual private gateway and attach it to your VPC. For more information about creating a virtual private gateway, see Adding a Hardware Virtual Private Gateway to Your VPC in the Amazon VPC User Guide.

To connect to public AWS products such as Amazon EC2 and Amazon S3, you need to provide the following:

A public ASN that you own (preferred) or a private ASN.

Public IP addresses (/31) (that is, one for each end of the BGP session) for each BGP session. If you do not have public IP addresses to assign to this connection, log on to AWS and then open a ticket with AWS Support.

The public routes that you will advertise over BGP.

NEW QUESTION 255

You are migrating a legacy client-server application to AWS. The application responds to a specific DNS domain (e.g. www.example.com) and has a 2-tier architecture, with multiple application servers and a database server. Remote clients use TCP to connect to the application servers. The application servers need to know the IP address of the clients in order to function properly and are currently taking that information from the TCP socket. A Multi-AZ RDS MySQL instance will be used for the database. During the migration you can change the application code, but you have to file a change request.

How would you implement the architecture on AWS in order to maximize scalability and high availability?

- A. File a change request to implement Alias Resource support in the application
- B. Use Route 53 Alias Resource Record to distribute load on two application servers in different AZs.
- C. File a change request to implement Latency Based Routing support in the application
- D. Use Route 53 with Latency Based Routing enabled to distribute load on two application servers in different AZs.
- E. File a change request to implement Cross-Zone support in the application
- F. Use an ELB with a TCP Listener and Cross-Zone Load Balancing enabled, two application servers in different AZs.
- G. File a change request to implement Proxy Protocol support in the application
- H. Use an ELB with a TCP Listener and Proxy Protocol enabled to distribute load on two application servers in different AZs.

Answer: D

NEW QUESTION 259

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

- A. Use S3 with reduced redundancy to store and serve the scanned files, install the commercial search application on EC2 Instances and configure with auto-scaling and an Elastic Load Balancer.
- B. Model the environment using CloudFormation use an EC2 instance running Apache webserver and an open source search application, stripe multiple standard EB5 volumes together to store the JPEGs and search index.
- C. Use S3 with standard redundancy to store and serve the scanned files, use CloudSearch for queryprocessing, and use Elastic Beanstalk to host the website across multiple availability zones.
- D. Use a single-AZ RD5 MySQL instance to store the search index and the JPEG images use an EC2 instance to serve the website and translate user queries into SQL.

E. Use a CloudFront download distribution to serve the JPEGs to the end users and Install the current commercial search product, along with a Java Container Tor the website on EC2 instances and use Route53 with DNS round-robin.

Answer: C

Explanation:

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

Storage Layer for 17TB of Images: This is the perfect use case for 53. Off-load all the web requests directly to the relevant JPEGs in 53. Your EC2 boxes just generate links to them.

If your app already serves it's own images (not links to images), you might start with EFS. But more than likely, you can just setup a web server to re-write or re-direct all JPEG links to 53 pretty easily.

If you use 53, don't serve directly from the bucket- Serve via a CNAME in domain you control. That way, you can switch in C|oudFront easily.

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

Cache Layer: Adding Cloud Front in front of 53 will help people on the other side of the world-- well, possibly. Typical archives follow a power law. The long tail of requests means that most JPEGs won't be requested enough to be in the cache. So you are only speeding up the most popular objects. You can always wait, and switch in CF later after you know your costs better. (In some cases, it can actually lower costs.)

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

Database Layer: A few options:

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

Use RDS to run MySQL/ Postgres

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

App Layer:

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

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

Preventing Abuse: When you put something in a public 53 bucket, people will hot-link it from their web pages. If you want to prevent that, your app on the EC2 box can generate signed links to 53 that expire in a few hours. Now everyone will be forced to go thru the app, and the app can apply rate limiting, etc. Saving money:

If you don't mind having downtime:

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

use Reduced Redundancy in 53 to save a few hundred bucks per month (Someone will have to "go fix it" every time it breaks, including having an off-line copy to repair 53.)

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

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

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

We're missing some information like load, latency expectations from search, indexing speed, size of the search index, etc. But with what you've given us, I would go with 53 as the storage for the files (53 rocks. It is really, really awesome). If you're stuck with the commercial search application, then on EC2 instances with autoscaling and an ELB. If you are allowed an alternative search engine, Elasticsearch is probably your best bet. I'd run it on EC2 instead ofthe AWS Elasticsearch service, as IMHO it's not ready yet. Don't autoscale Elasticsearch automatically though, it'll cause all sorts of issues. I have zero experience with CloudSearch so I can't comment on that. Regardless of which option, I'd use Cloud Formation for all of it.

NEW QUESTION 264

.....

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](#)