

AWS-Solution-Architect-Associate Dumps

AWS Certified Solutions Architect - Associate

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



NEW QUESTION 1

In Amazon EC2 Container Service components, what is the name of a logical grouping of container instances on which you can place tasks?

- A. A cluster
- B. A container instance
- C. A container
- D. A task definition

Answer: A

Explanation: Amazon ECS contains the following components:

A Cluster is a logical grouping of container instances that you can place tasks on.

A Container instance is an Amazon EC2 instance that is running the Amazon ECS agent and has been registered into a cluster.

A Task definition is a description of an application that contains one or more container definitions. A Scheduler is the method used for placing tasks on container instances.

A Service is an Amazon ECS service that allows you to run and maintain a specified number of instances of a task definition simultaneously.

A Task is an instantiation of a task definition that is running on a container instance. A Container is a Linux container that was created as part of a task.

Reference: <http://docs.aws.amazon.com/AmazonECS/latest/developerguide/Welcome.html>

NEW QUESTION 2

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 3

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

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

Answer: A

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

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

NEW QUESTION 4

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 5

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 local

- 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 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 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 AMI 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

What is a placement group in Amazon EC2?

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

Answer: A

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

NEW QUESTION 8

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 9

An edge location refers to which Amazon Web Service?

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

Answer: C

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

NEW QUESTION 10

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

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

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

Answer: D

Explanation: Amazon ElastiCache is a web service that makes it easy to deploy and run Memcached or Redis protocol-compliant server nodes in the cloud. Amazon ElastiCache improves the performance of web applications by allowing you to retrieve information from a fast, managed, in-memory caching system, instead of relying entirely on slower disk-based databases. The service simplifies and offloads the management, monitoring and operation of in-memory cache environments, enabling your engineering resources to focus on developing applications. Using Amazon ElastiCache, you can not only improve load and response times to user actions and queries, but also reduce the cost associated with scaling web applications.
Reference: <https://aws.amazon.com/elasticache/faqs/>

NEW QUESTION 10

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

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

Answer: D

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

NEW QUESTION 12

Your supervisor has asked you to build a simple file synchronization service for your department. He doesn't want to spend too much money and he wants to be notified of any changes to files by email. What do you think would be the best Amazon service to use for the email solution?

- A. Amazon SES
- B. Amazon CloudSearch
- C. Amazon SWF
- D. Amazon AppStream

Answer: A

Explanation: File change notifications can be sent via email to users following the resource with Amazon Simple Email Service (Amazon SES), an easy-to-use, cost-effective email solution.
Reference: http://media.amazonaws.com/architecturecenter/AWS_ac_ra_filesync_08.pdf

NEW QUESTION 17

Your manager has just given you access to multiple VPN connections that someone else has recently set up between all your company's offices. She needs you to make sure that the communication between the VPNs is secure. Which of the following services would be best for providing a low-cost hub-and-spoke model for

primary or backup connectivity between these remote offices?

- A. Amazon CloudFront
- B. AWS Direct Connect
- C. AWS CloudHSM
- D. AWS VPN CloudHub

Answer: D

Explanation: If you have multiple VPN connections, you can provide secure communication between sites using the AWS VPN CloudHub. The VPN CloudHub operates on a simple hub-and-spoke model that you can use with or without a VPC. This design is suitable for customers with multiple branch offices and existing Internet connections who would like to implement a convenient, potentially low-cost hub-and-spoke model for primary or backup connectivity between these remote offices.

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

NEW QUESTION 19

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 model.

Answer: C

Explanation: Amazon DynamoDB supports increment and decrement atomic operations.

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

NEW QUESTION 22

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

- A. 950
- B. 990
- C. 1000
- D. 900

Answer: D

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

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

NEW QUESTION 26

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

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

Answer: C

Explanation: AWS Import/Export supports: Import to Amazon S3

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

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

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

NEW QUESTION 28

In Amazon EC2, partial instance-hours are billed .

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

Answer: D

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

NEW QUESTION 30

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 31

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

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

Answer: B

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

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

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

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

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

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

NEW QUESTION 35

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

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

Answer: B

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

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

NEW QUESTION 37

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 41

Much of your company's data does not need to be accessed often, and can take several hours for retrieval time, so it's stored on Amazon Glacier. However

someone within your organization has expressed concerns that his data is more sensitive than the other data, and is wondering whether the high level of encryption that he knows is on S3 is also used on the much cheaper Glacier service. Which of the following statements would be most applicable in regards to this concern?

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

Answer: C

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

NEW QUESTION 46

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

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

Answer: A

Explanation: Several factors can affect the performance of Amazon EBS volumes, such as instance configuration, I/O characteristics, workload demand, and storage configuration. Frequent snapshots provide a higher level of data durability, but they may slightly degrade the performance of your application while the snapshot is in progress. This trade off becomes critical when you have data that changes rapidly. Whenever possible, plan for snapshots to occur during off-peak times in order to minimize workload impact.
Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSPerformance.html>

NEW QUESTION 48

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 52

A major finance organisation has engaged your company to set up a large data mining application. Using AWS you decide the best service for this is Amazon Elastic MapReduce(EMR) which you know uses Hadoop. Which of the following statements best describes Hadoop?

- A. Hadoop is 3rd Party software which can be installed using AMI
- B. Hadoop is an open source python web framework
- C. Hadoop is an open source Java software framework
- D. Hadoop is an open source javascript framework

Answer: C

Explanation: Amazon EMR uses Apache Hadoop as its distributed data processing engine. Hadoop is an open source, Java software framework that supports data-intensive distributed applications running on large clusters of commodity hardware. Hadoop implements a programming model named "MapReduce," where the data is dMded into many small fragments of work, each of which may be executed on any node in the cluster. This framework has been widely used by developers, enterprises and startups and has proven to be a reliable software platform for processing up to petabytes of data on clusters of thousands of commodity machines.
Reference: <http://aws.amazon.com/elasticmapreduce/faqs/>

NEW QUESTION 56

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 57

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

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

Answer: B

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

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

NEW QUESTION 62

You have just been given a scope for a new client who has an enormous amount of data (petabytes) that he constantly needs analysed. Currently he is paying a huge amount of money for a data warehousing company to do this for him and is wondering if AWS can provide a cheaper solution. Do you think AWS has a solution for this?

- A. Ye
- B. Amazon SimpleDB
- C. N
- D. Not presently
- E. Ye
- F. Amazon Redshift
- G. Ye
- H. Your choice of relational AMIs on Amazon EC2 and EBS

Answer: C

Explanation: Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse service that makes it simple and cost-effective to efficiently analyze all your data using your existing business intelligence tools. You can start small for just \$0.25 per hour with no commitments or upfront costs and scale to a petabyte or more for \$1,000 per terabyte per year, less than a tenth of most other data warehousing solutions. 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.

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

NEW QUESTION 63

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

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

Answer: C

Explanation: Snapshots with AWS Marketplace product codes can't be made public. Reference:

<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/ebs-modifying-snapshot-permissions.html>

NEW QUESTION 65

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 AWS

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 70

Which of the following AWS CLI commands is syntactically incorrect?

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

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

Answer: A

Explanation: The following CLI command is missing a hyphen before "-message".

```
aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError -message "Script Failure"
```

It has been added below in red

```
aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError ---message "Script Failure"
```

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

NEW QUESTION 75

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 80

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 83

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

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

Answer: B

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

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

All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.

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

NEW QUESTION 84

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

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

Answer: C

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

NEW QUESTION 86

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 88

What does the following policy for Amazon EC2 do?

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

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

Answer: A

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

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

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

NEW QUESTION 90

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 95

What does Amazon DynamoDB provide?

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

Answer: D

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

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

NEW QUESTION 100

You want to use AWS Import/Export to send data from your S3 bucket to several of your branch offices. What should you do if you want to send 10 storage units to AWS?

- A. Make sure your disks are encrypted prior to shipping.
- B. Make sure you format your disks prior to shipping.
- C. Make sure your disks are 1TB or more.
- D. Make sure you submit a separate job request for each device

Answer: D

Explanation: When using Amazon Import/Export, a separate job request needs to be submitted for each physical device even if they belong to the same import or export job.

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

NEW QUESTION 103

What would be the best way to retrieve the public IP address of your EC2 instance using the CLI?

- A. Using tags
- B. Using traceroute
- C. Using ipconfig
- D. Using instance metadata

Answer: D

Explanation: To determine your instance's public IP address from within the instance, you can use instance metadata. Use the following command to access the public IP address: For Linux use, \$ curl

<http://169.254.169.254/latest/meta-data/public-ipv4>, and for Windows use, \$ wget <http://169.254.169.254/latest/meta-data/public-ipv4>.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-instance-addressing.htm>

NEW QUESTION 106

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

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

Answer: D

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

storage configuration.

IOPS are input/output operations per second. Amazon EBS measures each I/O operation per second (that is 256 KB or smaller) as one IOPS. I/O operations that are larger than 256 KB are counted in 256 KB capacity units. For example, a 1,024 KB I/O operation would count as 4 IOPS.

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

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

NEW QUESTION 109

Having set up a website to automatically be redirected to a backup website if it fails, you realize that there are different types of failovers that are possible. You need all your resources to be available the majority of the time. Using Amazon Route 53 which configuration would best suit this requirement?

- A. Active-active failover.
- B. Non
- C. Route 53 can't failover.
- D. Active-passive failover.
- E. Active-active-passive and other mixed configuration

Answer: A

Explanation: You can set up a variety of failover configurations using Amazon Route 53 alias: weighted, latency, geolocation routing, and failover resource record sets.

Active-active failover: Use this failover configuration when you want all of your resources to be available the majority of the time. When a resource becomes unavailable, Amazon Route 53 can detect that it's unhealthy and stop including it when responding to queries.

Active-passive failover: Use this failover configuration when you want a primary group of resources to be available the majority of the time and you want a secondary group of resources to be on standby in case all of the primary resources become unavailable. When responding to queries, Amazon Route 53 includes only the healthy primary resources. If all of the primary resources are unhealthy, Amazon Route 53 begins to include only the healthy secondary resources in response to DNS queries.

Active-active-passive and other mixed configurations: You can combine alias and non-alias resource record sets to produce a variety of Amazon Route 53 behaviors.

Reference: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html>

NEW QUESTION 112

True or False: In Amazon Route 53, you can create a hosted zone for a top-level domain (TLD).

- A. FALSE
- B. False, Amazon Route 53 automatically creates it for you.
- C. True, only if you send an XML document with a CreateHostedZoneRequest element for TLD.
- D. TRUE

Answer: A

Explanation: In Amazon Route 53, you cannot create a hosted zone for a top-level domain (TLD).

Reference: http://docs.aws.amazon.com/Route53/latest/APIReference/API_CreateHostedZone.html

NEW QUESTION 116

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

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

Answer: C

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

The total volume of data and number of archives you can store are unlimited. Individual Amazon Glacier archives can range in size from 1 byte to 40 terabytes.

The largest archive that can be uploaded in a single upload request is 4 gigabytes.

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

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

NEW QUESTION 121

You are setting up your first Amazon Virtual Private Cloud (Amazon VPC) so you decide to use the VPC wizard in the AWS console to help make it easier for you. Which of the following statements is correct regarding instances that you launch into a default subnet via the VPC wizard?

- A. Instances that you launch into a default subnet receive a public IP address and 10 private IP addresses.
- B. Instances that you launch into a default subnet receive both a public IP address and a private IP address.
- C. Instances that you launch into a default subnet don't receive any IP addresses and you need to define them manually.
- D. Instances that you launch into a default subnet receive a public IP address and 5 private IP addresses.

Answer: B

Explanation: Instances that you launch into a default subnet receive both a public IP address and a private IP address. Instances in a default subnet also receive both public and private DNS hostnames. Instances that you launch into a nondefault subnet in a default VPC don't receive a public IP address or a DNS hostname. You can change your subnet's default public IP addressing behavior.

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

NEW QUESTION 125

Does AWS CloudFormation support Amazon EC2 tagging?

- A. Yes, AWS CloudFormation supports Amazon EC2 tagging
- B. No, CloudFormation doesn't support any tagging
- C. No, it doesn't support Amazon EC2 tagging.
- D. It depends if the Amazon EC2 tagging has been defined in the template

Answer: A

Explanation: In AWS CloudFormation, Amazon EC2 resources that support the tagging feature can also be tagged in an AWS template. The tag values can refer to template parameters, other resource names, resource attribute values (e.g. addresses), or values computed by simple functions (e.g., a concatenated list of strings).

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

NEW QUESTION 126

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

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

Answer: A

Explanation: By launching your instances into a VPC instead of EC2-Classic, you gain the ability to: Assign static private IP addresses to your instances that persist across starts and stops Assign multiple IP addresses to your instances

Define network interfaces, and attach one or more network interfaces to your instances Change security group membership for your instances while they're running

Control the outbound traffic from your instances (egress filtering) in addition to controlling the inbound traffic to them (ingress filtering)

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

Run your instances on single-tenant hardware

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

NEW QUESTION 128

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

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

Answer: B

Explanation: Amazon S3 supports several mechanisms that give you flexibility to control who can access your data as well as how, when, and where they can access it. Amazon S3 provides four different access control mechanisms: AWS Identity and Access Management (IAM) policies, Access Control Lists (ACLs), bucket policies, and query string authentication. IAM enables organizations to create and manage multiple users under a single AWS account. With IAM policies, you can grant IAM users fine-grained control to your Amazon S3 bucket or objects. You can use ACLs to selectively add (grant) certain permissions on individual objects.

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

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

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

NEW QUESTION 129

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

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

Answer: B

Explanation:

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

Reference:

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

NEW QUESTION 130

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

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

Answer: A

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

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

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

NEW QUESTION 132

You have set up an Elastic Load Balancer (ELB) with the usual default settings, which route each request independently to the application instance with the smallest load. However, someone has asked you to bind a user's session to a specific application instance so as to ensure that all requests coming from the user during the session will be sent to the same application instance. AWS has a feature to do this. What is it called?

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

Answer: D

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

An application instance must always receive and send two cookies: A cookie that defines the stickiness duration and a special Elastic Load Balancing cookie named AWSELB, that has the mapping to the application instance.

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

NEW QUESTION 133

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

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

Answer: C

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

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

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

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

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

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

NEW QUESTION 134

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

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

D. Create an Origin Access Identity (OAI) for CloudFront and grant access to the objects in your S3 bucket to that OAI.

Answer: D

Explanation: You restrict access to Amazon S3 content by creating an origin access identity, which is a special CloudFront user. You change Amazon S3 permissions to give the origin access identity permission to access your objects, and to remove permissions from everyone else. When your users access your Amazon S3 objects using CloudFront URLs, the CloudFront origin access identity gets the objects on your users' behalf. If your users try to access objects using Amazon S3 URLs, they're denied access. The origin access identity has permission to access objects in your Amazon S3 bucket, but users don't. Reference: <http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html>

NEW QUESTION 137

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

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

Answer: A

Explanation: The user can join multiple provisioned IOPS volumes together in a RAID 1 configuration to achieve better fault tolerance. RAID 1 does not provide a write performance improvement; it requires more bandwidth than non-RAID configurations since the data is written simultaneously to multiple volumes. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html>

NEW QUESTION 140

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 144

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 147

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 151

Which of the following strategies can be used to control access to your Amazon EC2 instances?

- A. DB security groups
- B. IAM policies
- C. None of these
- D. EC2 security groups

Answer: D

Explanation: IAM policies allow you to specify what actions your IAM users are allowed to perform against your EC2 Instances. However, when it comes to access control, security groups are what you need in order to define and control the way you want your instances to be accessed, and whether or not certain kind of communications are allowed or not.

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

NEW QUESTION 152

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

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

Answer: C

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

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

NEW QUESTION 155

In Amazon EC2, if your EBS volume stays in the detaching state, you can force the detachment by clicking .

- A. Force Detach
- B. Detach Instance
- C. AttachVolume
- D. AttachInstance

Answer: A

Explanation: If your volume stays in the detaching state, you can force the detachment by clicking Force Detach. Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/ebs-detaching-volume.html>

NEW QUESTION 159

Do you need to shutdown your EC2 instance when you create a snapshot of EBS volumes that serve as root devices?

- A. No, you only need to shutdown an instance before deleting it.
- B. Yes
- C. No, the snapshot would turn off your instance automatically.
- D. No

Answer: B

Explanation: Yes, to create a snapshot for Amazon EBS volumes that serve as root devices, you should stop the instance before taking the snapshot. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html>

NEW QUESTION 160

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 her

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 165

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 168

A user is running a batch process which runs for 1 hour every day. Which of the below mentioned options is the right instance type and costing model in this case if the user performs the same task for the whole year?

- A. EBS backed instance with on-demand instance pricing.
- B. EBS backed instance with heavy utilized reserved instance pricing.
- C. EBS backed instance with low utilized reserved instance pricing.
- D. Instance store backed instance with spot instance pricing.

Answer: A

Explanation: For Amazon Web Services, the reserved instance helps the user save money if the user is going to run the same instance for a longer period. Generally if the user uses the instances around 30-40% annually it is recommended to use RI. Here as the instance runs only for 1 hour daily it is not recommended to have RI as it will be costlier. The user should use on-demand with EBS in this case.

Reference: <http://aws.amazon.com/ec2/purchasing-options/reserved-instances/>

NEW QUESTION 173

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

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

Answer: A

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

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

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

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

NEW QUESTION 178

Which IAM role do you use to grant AWS Lambda permission to access a DynamoDB Stream?

- A. Dynamic role
- B. Invocation role
- C. Execution role
- D. Event Source role

Answer: C

Explanation: You grant AWS Lambda permission to access a DynamoDB Stream using an IAM role known as the "execution role".
Reference: <http://docs.aws.amazon.com/lambda/latest/dg/intro-permission-model.htm>

NEW QUESTION 182

Name the disk storage supported by Amazon Elastic Compute Cloud (EC2).

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

Answer: D

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

NEW QUESTION 187

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

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

Answer: A

Explanation: With IAM, you can centrally manage users, security credentials such as access keys, and permissions that control which AWS resources users can access.
In some cases, you might have an IAM user with full access to IAM and Amazon S3. If the IAM user assigns a bucket policy to an Amazon S3 bucket and doesn't specify the root user as a principal, the root user is denied access to that bucket. However, as the root user, you can still access the bucket by modifying the bucket policy to allow root user access.
Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/iam-troubleshooting.html#testing2>

NEW QUESTION 191

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

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

Answer: A

Explanation: Amazon CloudWatch provides the functionality to graph the metric data generated either by the AWS services or the custom metric to make it easier for the user to analyse. The AWS CloudWatch console provides the option to change the granularity of a graph and zoom in to see data over a shorter time period. To zoom, the user has to click in the graph details pane, drag on the graph area for selection, and then release the mouse button.
Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/zoom_in_on_graph.html

NEW QUESTION 192

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

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

Answer: A

Explanation: DynamoDB supports two types of secondary indexes:
Local secondary index — an index that has the same hash key as the table, but a different range key. A local secondary index is "local" in the sense that every partition of a local secondary index is scoped to a table partition that has the same hash key.
Global secondary index — an index with a hash and range key that can be different from those on the table. A global secondary index is considered "global" because queries on the index can span all of the data in a table, across all partitions.

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

NEW QUESTION 194

The common use cases for DynamoDB Fine-Grained Access Control (FGAC) are cases in which the end user wants .

- A. to change the hash keys of the table directly
- B. to check if an IAM policy requires the hash keys of the tables directly
- C. to read or modify any codecommit key of the table directly, without a middle-tier service
- D. to read or modify the table directly, without a middle-tier service

Answer: D

Explanation: FGAC can benefit any application that tracks information in a DynamoDB table, where the end user (or application client acting on behalf of an end user) wants to read or modify the table directly, without a middle-tier service. For instance, a developer of a mobile app named Acme can use FGAC to track the top score of every Acme user in a DynamoDB table. FGAC allows the application client to modify only the top score for the user that is currently running the application.

Reference: http://aws.amazon.com/dynamodb/faqs/#security_anchor

NEW QUESTION 197

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 198

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 203

A user comes to you and wants access to Amazon CloudWatch but only wants to monitor a specific LoadBalancer. Is it possible to give him access to a specific set of instances or a specific LoadBalancer?

- A. No because you can't use IAM to control access to CloudWatch data for specific resources.
- B. Yes
- C. You can use IAM to control access to CloudWatch data for specific resources.
- D. No because you need to be Sysadmin to access CloudWatch data.
- E. Yes
- F. Any user can see all CloudWatch data and needs no access right

Answer: A

Explanation: Amazon CloudWatch integrates with AWS Identity and Access Management (IAM) so that you can

specify which CloudWatch actions a user in your AWS Account can perform. For example, you could create an IAM policy that gives only certain users in your organization permission to use GetMetricStatistics. They could then use the action to retrieve data about your cloud resources.

You can't use IAM to control access to CloudWatch data for specific resources. For example, you can't give a user access to CloudWatch data for only a specific set of instances or a specific LoadBalancer. Permissions granted using IAM cover all the cloud resources you use with CloudWatch. In addition, you can't use IAM roles with the Amazon CloudWatch command line tools.

Using Amazon CloudWatch with IAM doesn't change how you use CloudWatch. There are no changes to CloudWatch actions, and no new CloudWatch actions related to users and access control.

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

NEW QUESTION 208

You have multiple VPN connections and want to provide secure communication between sites using the AWS VPN CloudHub. Which statement is the most accurate in describing what you must do to set this up correctly?

- A. Create a virtual private gateway with multiple customer gateways, each with unique Border Gateway Protocol (BGP) Autonomous System Numbers (ASNs)
- B. Create a virtual private gateway with multiple customer gateways, each with a unique set of keys
- C. Create a virtual public gateway with multiple customer gateways, each with a unique Private subnet
- D. Create a virtual private gateway with multiple customer gateways, each with unique subnet id

Answer: A

Explanation: If you have multiple VPN connections, you can provide secure communication between sites using the AWS VPN CloudHub. The VPN CloudHub operates on a simple hub-and-spoke model that you can use with or without a VPC. This design is suitable for customers with multiple branch offices and existing Internet connections who'd like to implement a convenient, potentially low-cost hub-and-spoke model for primary or backup connectivity between these remote offices.

To use the AWS VPN CloudHub, you must create a virtual private gateway with multiple customer gateways, each with unique Border Gateway Protocol (BGP) Autonomous System Numbers (ASNs). Customer gateways advertise the appropriate routes (BGP prefixes) over their VPN connections. These routing advertisements are received and re-advertised to each BGP peer, enabling each site to send data to and receive data from the other sites. The routes for each spoke must have unique ASNs and the sites must not have overlapping IP ranges. Each site can also send and receive data from the VPC as if they were using a standard VPN connection.

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

NEW QUESTION 211

Which one of the below is not an AWS Storage Service?

- A. Amazon S3
- B. Amazon Glacier
- C. Amazon CloudFront
- D. Amazon EBS

Answer: C

Explanation: AWS Storage Services are: Amazon S3
Amazon Glacier
Amazon EBS
AWS Storage Gateway

Reference: <https://console.aws.amazon.com/console>

NEW QUESTION 215

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

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

Answer: A

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

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

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

NEW QUESTION 217

After setting up an EC2 security group with a cluster of 20 EC2 instances, you find an error in the security group settings. You quickly make changes to the security group settings. When will the changes to the settings be effective?

- A. The settings will be effective immediately for all the instances in the security group.
- B. The settings will be effective only when all the instances are restarted.
- C. The settings will be effective for all the instances only after 30 minutes.
- D. The settings will be effective only for the new instances added to the security group.

Answer: A

Explanation: Amazon Redshift applies changes to a cluster security group immediately. So if you have associated the cluster security group with a cluster, inbound cluster access rules in the updated cluster security group apply immediately.

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

NEW QUESTION 218

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

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

Answer: B

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

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

NEW QUESTION 220

You are setting up some EBS volumes for a customer who has requested a setup which includes a RAID (redundant array of inexpensive disks). AWS has some recommendations for RAID setups. Which RAID setup is not recommended for Amazon EBS?

- A. RAID 5 only
- B. RAID 5 and RAID 6
- C. RAID 1 only
- D. RAID 1 and RAID 6

Answer: B

Explanation: With Amazon EBS, you can use any of the standard RAID configurations that you can use with a traditional bare metal server, as long as that particular RAID configuration is supported by the operating system for your instance. This is because all RAID is accomplished at the software level. For greater I/O performance than you can achieve with a single volume, RAID 0 can stripe multiple volumes together; for on-instance redundancy, RAID 1 can mirror two volumes together.

RAID 5 and RAID 6 are not recommended for Amazon EBS because the parity write operations of these RAID modes consume some of the IOPS available to your volumes.

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

NEW QUESTION 223

You receive the following request from a client to quickly deploy a static website for them, specifically on AWS. The requirements are low-cost, reliable, online storage, and a reliable and cost-effective way to route customers to the website, as well as a way to deliver content with low latency and high data transfer speeds so that visitors to his website don't experience unnecessary delays. What do you think would be the minimum AWS services that could fulfill the client's request?

- A. Amazon Route 53, Amazon CloudFront and Amazon VPC.
- B. Amazon S3, Amazon Route 53 and Amazon RDS
- C. Amazon S3, Amazon Route 53 and Amazon CloudFront
- D. Amazon S3 and Amazon Route 53.

Answer: C

Explanation: You can easily and inexpensively use AWS to host a website that uses client-side technologies (such as HTML, CSS, and JavaScript) and does not require server-side technologies (such as PHP and ASP.NET). This type of site is called a static website, and is used to display content that does not change frequently. Before you create and deploy a static website, you must plan your architecture to ensure that it meets your requirements. Amazon S3, Amazon Route 53, and Amazon CloudFront would be required in this instance.

Reference: <http://docs.aws.amazon.com/gettingstarted/latest/swf/website-hosting-intro.html>

NEW QUESTION 228

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 229

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.xlarge instance is 10 Gigabit. Reference: <http://aws.amazon.com/ec2/instance-types/>

NEW QUESTION 234

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

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

Answer: B

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

NEW QUESTION 235

You need to create a JSON-formatted text file for AWS CloudFormation. This is your first template and the only thing you know is that the templates include several major sections but there is only one that is required for it to work. What is the only section required?

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

Answer: C

Explanation: AWS CloudFormation is a service that helps you model and set up your Amazon Web Services resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS. You create a template that describes all the AWS resources that you want (like Amazon EC2 instances or Amazon RDS DB instances), and AWS CloudFormation takes care of provisioning and configuring those resources for you.

A template is a JSON-formatted text file that describes your AWS infrastructure. Templates include several major sections.

The Resources section is the only section that is required.

The first character in the template must be an open brace ({), and the last character must be a closed brace (}). The following template fragment shows the template structure and sections.

Reference: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-anatomy.html>

NEW QUESTION 238

How long does an AWS free usage tier EC2 last for?

- A. Forever
- B. 12 Months upon signup
- C. 1 Month upon signup
- D. 6 Months upon signup

Answer: B

Explanation: The AWS free usage tier will expire 12 months from the date you sign up. When your free usage expires or if your application use exceeds the free usage tiers, you simply pay the standard, pay-as-you-go service rates.

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

NEW QUESTION 240

A user is hosting a website in the US West-1 region. The website has the highest client base from the Asia-Pacific (Singapore / Japan) region. The application is accessing data from S3 before serving it to client. Which of the below mentioned regions gives a better performance for S3 objects?

- A. Japan
- B. Singapore
- C. US East
- D. US West-1

Answer: D

Explanation: Access to Amazon S3 from within Amazon EC2 in the same region is fast. In this aspect, though the client base is Singapore, the application is being hosted in the US West-1 region. Thus, it is recommended that S3 objects be stored in the US-West-1 region.

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

NEW QUESTION 242

You have been setting up an Amazon Virtual Private Cloud (Amazon VPC) for your company, including setting up subnets. Security is a concern, and you are not sure which is the best security practice for securing subnets in your VPC. Which statement below is correct in describing the protection of AWS resources in each subnet?

- A. You can use multiple layers of security, including security groups and network access control lists (ACL).
- B. You can only use access control lists (ACL).

- C. You don't need any security in subnets.
- D. You can use multiple layers of security, including security groups, network access control lists (ACL) and CloudHSM.

Answer: A

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

To protect the AWS resources in each subnet, you can use multiple layers of security, including security groups and network access control lists (ACL).

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

NEW QUESTION 243

Your organization is in the business of architecting complex transactional databases. For a variety of reasons, this has been done on EBS. What is AWS's recommendation for customers who have architected databases using EBS for backups?

- A. Backups to Amazon S3 be performed through the database management system.
- B. Backups to AWS Storage Gateway be performed through the database management system.
- C. If you take regular snapshots no further backups are required.
- D. Backups to Amazon Glacier be performed through the database management system.

Answer: A

Explanation: Data stored in Amazon EBS volumes is redundantly stored in multiple physical locations as part of normal operation of those services and at no additional charge.

However, Amazon EBS replication is stored within the same availability zone, not across multiple zones; therefore, it is highly recommended that you conduct regular snapshots to Amazon S3 for long-term data durability.

For customers who have architected complex transactional databases using EBS, it is recommended that backups to Amazon S3 be performed through the database management system so that distributed transactions and logs can be checkpointed.

AWS does not perform backups of data that are maintained on virtual disks attached to running instances on Amazon EC2.

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

NEW QUESTION 248

You have three Amazon EC2 instances with Elastic IP addresses in the US East (Virginia) region, and you want to distribute requests across all three IPs evenly for users for whom US East (Virginia) is the appropriate region.

How many EC2 instances would be sufficient to distribute requests in other regions?

- A. 3
- B. 9
- C. 2
- D. 1

Answer: D

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

For example, suppose you have three Amazon EC2 instances with Elastic IP addresses in the US East (Virginia) region and you want to distribute requests across all three IPs evenly for users for whom US East (Virginia) is the appropriate region. Just one Amazon EC2 instance is sufficient in the other regions, although you can apply the same technique to many regions at once.

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

NEW QUESTION 250

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 253

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 258

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

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

Answer: B

Explanation: Amazon Elastic Map Reduce (EMR) is a web service that enables businesses, researchers, data analysts, and developers to easily and cost-effectively process vast amounts of data. Amazon EMR historically referred to an Amazon EMR cluster (and all processing steps assigned to it) as a "cluster". Every cluster has a unique identifier that starts with "j-". The different cluster states of an Amazon EMR cluster are listed below. STARTING — The cluster provisions, starts, and configures EC2 instances. BOOTSTRAPPING — Bootstrap actions are being executed on the cluster. RUNNING — A step for the cluster is currently being run. WAITING — The cluster is currently active, but has no steps to run. TERMINATING - The cluster is in the process of shutting down. TERMINATED - The cluster was shut down without error. TERMINATED_WITH_ERRORS - The cluster was shut down with errors.
Reference: <https://aws.amazon.com/elasticmapreduce/faqs/>

NEW QUESTION 259

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

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

Answer: B

Explanation: The AWS CloudHSM service defines a resource known as a high-availability (HA) partition group, which is a virtual partition that represents a group of partitions, typically distributed between several physical HSMs for high-availability.
Reference: <http://docs.aws.amazon.com/cloudhsm/latest/userguide/configuring-ha.html>

NEW QUESTION 261

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 264

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 265

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 269

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

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

Answer: D

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

If transmission of any part fails, you can re-transmit that part without affecting other parts. After all parts of your object are uploaded, Amazon S3 assembles these parts and creates the object. In general, when

your object size reaches 100 MB, you should consider using multipart uploads instead of uploading the object in a single operation.

Using multipart upload provides the following advantages:

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

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

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

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

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

NEW QUESTION 273

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 274

What happens to Amazon EBS root device volumes, by default, when an instance terminates?

- A. Amazon EBS root device volumes are moved to IAM.
- B. Amazon EBS root device volumes are copied into Amazon RDS.
- C. Amazon EBS root device volumes are automatically deleted.

D. Amazon EBS root device volumes remain in the database until you delete the

Answer: C

Explanation: By default, Amazon EBS root device volumes are automatically deleted when the instance terminates. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/terminating-instances.html>

NEW QUESTION 276

A gaming company comes to you and asks you to build them infrastructure for their site. They are not sure how big they will be as with all start ups they have limited money and big ideas. What they do tell you is that if the game becomes successful, like one of their previous games, it may rapidly grow to millions of users and generate tens (or even hundreds) of thousands of writes and reads per second. After considering all of this, you decide that they need a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. Which of the following databases do you think would best fit their needs?

- A. Amazon DynamoDB
- B. Amazon Redshift
- C. Any non-relational database.
- D. Amazon SimpleDB

Answer: A

Explanation: Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. Amazon DynamoDB enables customers to offload the administrative burdens of operating and scaling distributed databases to AWS, so they don't have to worry about hardware provisioning, setup and configuration, replication, software patching, or cluster scaling. Today's web-based applications generate and consume massive amounts of data. For example, an online game might start out with only a few thousand users and a light database workload consisting of 10 writes per second and 50 reads per second. However, if the game becomes successful, it may rapidly grow to millions of users and generate tens (or even hundreds) of thousands of writes and reads per second. It may also create terabytes or more of data per day. Developing your applications against Amazon DynamoDB enables you to start small and simply dial-up your request capacity for a table as your requirements scale, without incurring downtime. You pay highly cost-efficient rates for the request capacity you provision, and let Amazon DynamoDB do the work over partitioning your data and traffic over sufficient server capacity to meet your needs. Amazon DynamoDB does the database management and administration, and you simply store and request your data. Automatic replication and failover provides built-in fault tolerance, high availability, and data durability. Amazon DynamoDB gives you the peace of mind that your database is fully managed and can grow with your application requirements. Reference: <http://aws.amazon.com/dynamodb/faqs/>

NEW QUESTION 281

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 282

You're trying to delete an SSL certificate from the IAM certificate store, and you're getting the message "Certificate: <certificate-id> is being used by CloudFront." Which of the following statements is probably the reason why you are getting this error?

- A. Before you can delete an SSL certificate, you need to either rotate SSL certificates or revert from using a custom SSL certificate to using the default CloudFront certificate.
- B. You can't delete SSL certificates. You need to request it from AWS.
- C. Before you can delete an SSL certificate, you need to set up the appropriate access level in IAM
- D. Before you can delete an SSL certificate you need to set up https on your server

Answer: A

Explanation: CloudFront is a web service that speeds up distribution of your static and dynamic web content, for example, .html, .css, .php, and image files, to end users.

Every CloudFront web distribution must be associated either with the default CloudFront certificate or with a custom SSL certificate. Before you can delete an SSL certificate, you need to either rotate SSL certificates (replace the current custom SSL certificate with another custom SSL certificate) or revert from using a custom SSL certificate to using the default CloudFront certificate.

Reference: <http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Troubleshooting.htm>

NEW QUESTION 285

How many types of block devices does Amazon EC2 support?

- A. 4
- B. 5
- C. 2
- D. 1

Answer: C

Explanation: Amazon EC2 supports 2 types of block devices. Reference:
<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/block-device-mapping-concepts.html>

NEW QUESTION 290

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

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

Answer: C

Explanation: Amazon VPC provides two features that you can use to increase security for your VPC:
Security groups—Act as a firewall for associated Amazon EC2 instances, controlling both inbound and outbound traffic at the instance level and supports allow rules only.
Network access control lists (ACLs)—Act as a firewall for associated subnets, controlling both inbound and outbound traffic at the subnet level and supports allow rules and deny rules.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Security.html

NEW QUESTION 293

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 298

A user wants to increase the durability and availability of the EBS volume. Which of the below mentioned actions should he perform?

- A. Take regular snapshots.
- B. Create an AMI.
- C. Create EBS with higher capacity.
- D. Access EBS regularly

Answer: A

Explanation: In Amazon Web Services, Amazon EBS volumes that operate with 20 GB or less of modified data since their most recent snapshot can expect an annual failure rate (AFR) between 0.1% and 0.5%. For this reason, to maximize both durability and availability of their Amazon EBS data, the user should frequently create snapshots of the Amazon EBS volumes.
Reference: http://media.amazonwebservices.com/AWS_Storage_Options.pdf

NEW QUESTION 303

In relation to AWS CloudHSM, High-availability (HA) recovery is hands-off resumption by failed HA group members. Prior to the introduction of this function, the HA feature provided redundancy and performance, but required that a failed/lost group member be reinstated.

- A. automatically
- B. periodically
- C. manually
- D. continuously

Answer: C

Explanation: In relation to AWS CloudHSM, High-availability (HA) recovery is hands-off resumption by failed HA group members. Prior to the introduction of this function, the HA feature provided redundancy and performance, but required that a failed/lost group member be manually reinstated.
Reference: <http://docs.aws.amazon.com/cloudhsm/latest/userguide/ha-best-practices.html>

NEW QUESTION 308

Any person or application that interacts with AWS requires security credentials. AWS uses these credentials to identify who is making the call and whether to allow the requested access. You have just set up a VPC network for a client and you are now thinking about the best way to secure this network. You set up a security group called vpcsecuritygroup. Which following statement is true in respect to the initial settings that will be applied to this security group if you choose to use the default settings for this group?

- A. Allow all inbound traffic and allow no outbound traffic.
- B. Allow no inbound traffic and allow all outbound traffic.
- C. Allow inbound traffic on port 80 only and allow all outbound traffic.
- D. Allow all inbound traffic and allow all outbound traffic.

Answer: B

Explanation: Amazon VPC provides advanced security features such as security groups and network access control lists to enable inbound and outbound filtering at the instance level and subnet level.

AWS assigns each security group a unique ID in the form sg-xxxxxxx. The following are the initial settings for a security group that you create:

Allow no inbound traffic Allow all outbound traffic

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

NEW QUESTION 313

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

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

Answer: D

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

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

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

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

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

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

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

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

NEW QUESTION 318

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 323

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 328

Amazon Elastic Load Balancing is used to manage traffic on a fleet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits. Which of the following is not an advantage of ELB over an on-premise load balancer?

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

Answer: A

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

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

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

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

Supports end-to-end traffic encryption using TLS (previously SSL) on those networks that use secure HTTP (HTTPS) connections. When TLS is used, the TLS server certificate used to terminate client connections can be managed centrally on the load balancer, rather than on every individual instance. Reference: <http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf>

NEW QUESTION 330

An organization has a statutory requirement to protect the data at rest for the S3 objects. Which of the below mentioned options need not be enabled by the organization to achieve data security?

- A. MFA delete for S3 objects
- B. Client side encryption
- C. Bucket versioning
- D. Data replication

Answer: D

Explanation: AWS S3 provides multiple options to achieve the protection of data at REST. The options include Permission (Policy), Encryption (Client and Server Side), Bucket Versioning and MFA based delete. The user can enable any of these options to achieve data protection. Data replication is an internal facility by AWS where S3 replicates each object across all the Availability Zones and the organization need not enable it in this case.

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

NEW QUESTION 335

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 336

Which of the following features are provided by Amazon EC2?

- A. Exadata Database Machine, Optimized Storage Management, Flashback Technology, and Data Warehousing
- B. Instances, Amazon Machine Images (AMIs), Key Pairs, Amazon EBS Volumes, Firewall, Elastic IP address, Tags, and Virtual Private Clouds (VPCs)
- C. Real Application Clusters (RAC), ElastiCache Machine Images (EMIs), Data Warehousing, Flashback Technology, Dynamic IP address
- D. Exadata Database Machine, Real Application Clusters (RAC), Data Guard, Table and Index Partitioning, and Data Pump Compression

Answer: B

Explanation: Amazon EC2 provides the following features:

- Virtual computing environments, known as instances;
- Pre-configured templates for your instances, known as Amazon Machine Images (AMIs), that package the bits you need for your server (including the operating system and additional software)
- Various configurations of CPU, memory, storage, and networking capacity for your instances, known as instance types
- Secure login information for your instances using key pairs (AWS stores the public key, and you store the private key in a secure place)
- Storage volumes for temporary data that's deleted when you stop or terminate your instance, known as instance store volumes
- Persistent storage volumes for your data using Amazon Elastic Block Store (Amazon EBS), known as Amazon EBS volumes

- Multiple physical locations for your resources, such as instances and Amazon EBS volumes, known as regions and Availability Zones
- A firewall that enables you to specify the protocols, ports, and source IP ranges that can reach your instances using security groups
- Static IP addresses for dynamic cloud computing, known as Elastic IP addresses
- Metadata, known as tags, that you can create and assign to your Amazon EC2 resources
- Virtual networks you can create that are logically isolated from the rest of the AWS cloud, and that you can optionally connect to your own network, known as virtual private clouds (VPCs).

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

NEW QUESTION 339

A friend tells you he is being charged \$100 a month to host his WordPress website, and you tell him you can move it to AWS for him and he will only pay a fraction of that, which makes him very happy. He then tells you he is being charged \$50 a month for the domain, which is registered with the same people that set it up, and he asks if it's possible to move that to AWS as well. You tell him you aren't sure, but will look into it. Which of the following statements is true in regards to transferring domain names to AWS?

- A. You can't transfer existing domains to AWS.
- B. You can transfer existing domains into Amazon Route 53's management.
- C. You can transfer existing domains via AWS Direct Connect.
- D. You can transfer existing domains via AWS Import/Export

Answer: B

Explanation: With Amazon Route 53, you can create and manage your public DNS records with the AWS Management Console or with an easy-to-use API. If you need a domain name, you can find an available name and register it using Amazon Route 53. You can also transfer existing domains into Amazon Route 53's management.

Reference: <http://aws.amazon.com/route53/>

NEW QUESTION 341

In AWS CloudHSM, in addition to the AWS recommendation that you use two or more HSM appliances in a high-availability configuration to prevent the loss of keys and data, you can also perform a remote backup/restore of a Luna SA partition if you have purchased a:

- A. Luna Restore HSM.
- B. Luna Backup HSM.
- C. Luna HSNI.
- D. Luna SA HSM.

Answer: B

Explanation: In AWS CloudHSM, you can perform a remote backup/restore of a Luna SA partition if you have purchased a Luna Backup HSM.

Reference: <http://docs.aws.amazon.com/cloudhsm/latest/userguide/cloud-hsm-backup-restore.html>

NEW QUESTION 344

A user has launched a large EBS backed EC2 instance in the US-East-1a region. The user wants to achieve Disaster Recovery (DR) for that instance by creating another small instance in Europe. How can the user achieve DR?

- A. Copy the instance from the US East region to the EU region
- B. Use the "Launch more like this" option to copy the instance from one region to another
- C. Copy the running instance using the "Instance Copy" command to the EU region
- D. Create an AMI of the instance and copy the AMI to the EU region
- E. Then launch the instance from the EU AMI

Answer: D

Explanation: To launch an EC2 instance it is required to have an AMI in that region. If the AMI is not available in that region, then create a new AMI or use the copy command to copy the AMI from one region to the other region.

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

NEW QUESTION 347

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

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

Answer: C

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

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

NEW QUESTION 352

Content and IV|edia Server is the latest requirement that you need to meet for a client.

The client has been very specific about his requirements such as low latency, high availability, durability, and access control. Potentially there will be millions of views on this server and because of "spiky" usage patterns, operations teams will need to provision static hardware, network, and management resources to support the maximum expected need. The Customer base will be initially low but is expected to grow and become more geographically distributed.

Which of the following would be a good solution for content distribution?

- A. Amazon S3 as both the origin server and for caching
- B. AWS Storage Gateway as the origin server and Amazon EC2 for caching
- C. AWS CloudFront as both the origin server and for caching
- D. Amazon S3 as the origin server and Amazon CloudFront for caching

Answer: D

Explanation: As your customer base grows and becomes more geographically distributed, using a high- performance edge cache like Amazon CloudFront can provide substantial improvements in latency, fault tolerance, and cost.

By using Amazon S3 as the origin server for the Amazon CloudFront distribution, you gain the advantages of fast in-network data transfer rates, simple publishing/caching workflow, and a unified security framework.

Amazon S3 and Amazon CloudFront can be configured by a web service, the AWS Management Console, or a host of third-party management tools.

Reference:http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_media_02.pdf

NEW QUESTION 357

You are setting up your first Amazon Virtual Private Cloud (Amazon VPC) network so you decide you should probably use the AWS Management Console and the VPC Wizard. Which of the following is not an option for network architectures after launching the "Start VPC Wizard" in Amazon VPC page on the AWS Management Console?

- A. VPC with a Single Public Subnet Only
- B. VPC with a Public Subnet Only and Hardware VPN Access
- C. VPC with Public and Private Subnets and Hardware VPN Access
- D. VPC with a Private Subnet Only and Hardware VPN Access

Answer: B

Explanation: Amazon VPC enables you to build a virtual network in the AWS cloud - no VPNs, hardware, or physical datacenters required.

Your AWS resources are automatically provisioned in a ready-to-use default VPC. You can choose to create additional VPCs by going to Amazon VPC page on the AWS Management Console and click on the "Start VPC Wizard" button.

You'll be presented with four basic options for network architectures. After selecting an option, you can modify the size and IP address range of the VPC and its subnets. If you select an option with Hardware VPN Access, you will need to specify the IP address of the VPN hardware on your network. You can modify the VPC to add more subnets or add or remove gateways at any time after the VPC has been created.

The four options are:

VPC with a Single Public Subnet Only VPC with Public and Private Subnets

VPC with Public and Private Subnets and Hardware VPN Access VPC with a Private Subnet Only and Hardware VPN Access Reference:

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

NEW QUESTION 362

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 365

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 368

You are architecting a highly-scalable and reliable web application which will have a huge amount of content. You have decided to use CloudFront as you know it will speed up distribution of your static and dynamic web content and know that Amazon CloudFront integrates with Amazon CloudWatch metrics so that you can monitor your web application. Because you live in Sydney you have chosen the the Asia Pacific (Sydney) region in the AWS console. However you have set up this up but no CloudFront metrics seem to be appearing in the CloudWatch console. What is the most likely reason from the possible choices below for this?

- A. Metrics for CloudWatch are available only when you choose the same region as the application you are monitoring.
- B. You need to pay for CloudWatch for it to become active.
- C. Metrics for CloudWatch are available only when you choose the US East (Virginia)
- D. Virginia)
- E. Metrics for CloudWatch are not available for the Asia Pacific region as yet

Answer: C

Explanation: CloudFront is a global service, and metrics are available only when you choose the US East (N. Virginia) region in the AWS console. If you choose another region, no CloudFront metrics will appear in the CloudWatch console.

Reference:

<http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/monitoring-using-cloudwatch.html>

NEW QUESTION 370

After a major security breach your manager has requested a report of all users and their credentials in AWS. You discover that in IAM you can generate and download a credential report that lists all users in your account and the status of their various credentials, including passwords, access keys, MFA devices, and signing certificates. Which following statement is incorrect in regards to the use of credential reports?

- A. Credential reports are downloaded XML files.
- B. You can get a credential report using the AWS Management Console, the AWS CLI, or the IAM API.
- C. You can use the report to audit the effects of credential lifecycle requirements, such as password rotation.
- D. You can generate a credential report as often as once every four hour

Answer: A

Explanation: To access your AWS account resources, users must have credentials.

You can generate and download a credential report that lists all users in your account and the status of their various credentials, including passwords, access keys, MFA devices, and signing certificates. You can get a credential report using the AWS Management Console, the AWS CLI, or the IAM API.

You can use credential reports to assist in your auditing and compliance efforts. You can use the report to audit the effects of credential lifecycle requirements, such as password rotation. You can provide the report to an external auditor, or grant permissions to an auditor so that he or she can download the report directly.

You can generate a credential report as often as once every four hours. When you request a report, IAM first checks whether a report for the account has been generated within the past four hours. If so, the most recent report is downloaded. If the most recent report for the account is more than four hours old, or if there are no previous reports for the account, IAM generates and downloads a new report.

Credential reports are downloaded as comma-separated values (CSV) files.

You can open CSV files with common spreadsheet software to perform analysis, or you can build an application that consumes the CSV files programmatically and performs custom analysis. Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/credential-reports.html>

NEW QUESTION 374

In the most recent company meeting, your CEO focused on the fact that everyone in the organization needs to make sure that all of the infrastructure that is built is truly scalable. Which of the following statements is incorrect in reference to scalable architecture?

- A. A scalable service is capable of handling heterogeneity.
- B. A scalable service is resilient.
- C. A scalable architecture won't be cost effective as it grows.
- D. Increasing resources results in a proportional increase in performance

Answer: C

Explanation: In AWS it is critical to build a scalable architecture in order to take advantage of a scalable infrastructure. The cloud is designed to provide conceptually infinite scalability. However, you cannot leverage all that scalability in infrastructure if your architecture is not scalable. Both have to work together.

You will have to identify the monolithic components and bottlenecks in your architecture, identify the areas where you cannot leverage the on-demand provisioning capabilities in your architecture, and work to refactor your application, in order to leverage the scalable infrastructure and take advantage of the cloud.

Characteristics of a truly scalable application:

Increasing resources results in a proportional increase in performance A scalable service is capable of handling heterogeneity

A scalable service is operationally efficient A scalable service is resilient

A scalable service should become more cost effective when it grows (Cost per unit reduces as the number of units increases)

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

NEW QUESTION 378

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 380

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 385

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 388

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

- A. You can change the outbound rules for EC2-Classic
- B. Also, you can add and remove rules to a group at any time.
- C. You can modify an existing rule in a group
- 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-Classic
- G. However, you can add and remove rules to a group at any time

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 391

Which DNS name can only be resolved within Amazon EC2?

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

Answer: B

Explanation: Only Internal DNS name can be resolved within Amazon EC2. Reference:

<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-instance-addressing.html>

NEW QUESTION 394

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

- A. You can detach secondary (ethN) network interfaces when the instance is running or stopped
- B. However, you can't detach the primary (eth0) interface.
- C. Launching an instance with multiple network interfaces automatically configures interfaces, private IP addresses, and route tables on the operating system of

the instance.

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

E. Attaching another network interface to an instance is a valid method to increase or double the network bandwidth to or from the dual-homed instance

Answer: D

Explanation: Best Practices for Configuring Network Interfaces

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

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

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

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

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

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

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

Reference:

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

NEW QUESTION 399

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

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

Answer: D

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

Reference:

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

NEW QUESTION 402

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

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

Answer: D

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

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

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

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

Simplicity

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

NEW QUESTION 406

A user has launched an EC2 instance. The instance got terminated as soon as it was launched. Which of the below mentioned options is not a possible reason for this?

- A. The user account has reached the maximum volume limit
- B. The AMI is missin
- C. It is the required part
- D. The snapshot is corrupt
- E. The user account has reached the maximum EC2 instance limit

Answer: D

Explanation: When the user account has reached the maximum number of EC2 instances, it will not be allowed to launch an instance. AWS will throw an 'Instance Limit Exceeded' error. For all other reasons, such as

"AMI is missing part", "Corrupt Snapshot" or "Volume limit has reached" it will launch an EC2 instance and then terminate it.

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

NEW QUESTION 407

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

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

Answer: B

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

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

NEW QUESTION 410

You can seamlessly join an EC2 instance to your directory domain. What connectMty do you need to be able to connect remotely to this instance?

- A. You must have IP connectMty to the instance from the network you are connecting from.
- B. You must have the correct encryption keys to connect to the instance remotely.
- C. You must have enough bandwidth to connect to the instance.
- D. You must use MFA authentication to be able to connect to the instance remotel

Answer: A

Explanation: You can seamlessly join an EC2 instance to your directory domain when the instance is launched using the Amazon EC2 Simple Systems Manager. If you need to manually join an EC2 instance to your domain, you must launch the instance in the proper region and security group or subnet, then join the instance to the domain. To be able to connect remotely to these instances, you must have IP connectMty to the instances from the network you are connecting from. In most cases, this requires that an Internet gateway be attached to your VPC and that the instance has a public IP address.

Reference: http://docs.aws.amazon.com/directoryservice/latest/admin-guide/join_a_directory.html

NEW QUESTION 412

George has launched three EC2 instances inside the US-East-1a zone with his AWS account. Ray has launched two EC2 instances in the US-East-1a zone with his AWS account. Which of the below mentioned statements will help George and Ray understand the availability zone (AZ) concept better?

- A. All the instances of George and Ray can communicate over a private IP with a minimal cost
- B. The US-East-1a region of George and Ray can be different availability zones
- C. All the instances of George and Ray can communicate over a private IP without any cost
- D. The instances of George and Ray will be running in the same data centre

Answer: B

Explanation: Each AWS region has multiple, isolated locations known as Availability Zones. To ensure that the AWS resources are distributed across the Availability Zones for a region, AWS independently maps the Availability Zones to identifiers for each account. In this case the Availability Zone US-East-1a where George's EC2 instances are running might not be the same location as the US-East-1a zone of Ray's EC2 instances. There is no way for the user to coordinate the Availability Zones between accounts.

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

NEW QUESTION 417

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 419

A user has created an ELB with Auto Scaling. Which of the below mentioned offerings from ELB helps the user to stop sending new requests traffic from the load balancer to the EC2 instance when the instance is being deregistered while continuing in-flight requests?

- A. ELB sticky session
- B. ELB deregistration check
- C. ELB auto registration Off
- D. ELB connection draining

Answer: D

Explanation: The Elastic Load Balancer connection draining feature causes the load balancer to stop sending new requests to the back-end instances when the instances are deregistering or become unhealthy, while ensuring that in-flight requests continue to be served.

Reference:

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

NEW QUESTION 421

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

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

Answer: C

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

RRS is ideal for non-critical or reproducible data.

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

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

NEW QUESTION 422

A user is making a scalable web application with compartmentalization. The user wants the log module to be able to be accessed by all the application functionalities in an asynchronous way. Each module of the application sends data to the log module, and based on the resource availability it will process the logs. Which AWS service helps this functionality?

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

Answer: A

Explanation: Amazon Simple Queue Service (SQS) is a highly reliable distributed messaging system for storing messages as they travel between computers. By using Amazon SQS, developers can simply move data between distributed application components. It is used to achieve compartmentalization or loose coupling. In this case all the modules will send a message to the logger queue and the data will be processed by queue as per the resource availability.

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

NEW QUESTION 425

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

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

Answer: A

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

Amazon S3 provides four different access control mechanisms:

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

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

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

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

NEW QUESTION 429

Which one of the following can't be used as an origin server with Amazon CloudFront?

- A. A web server running in your infrastructure
- B. Amazon S3
- C. Amazon Glacier
- D. A web server running on Amazon EC2 instances

Answer: C

Explanation: Amazon CloudFront is designed to work with Amazon S3 as your origin server, customers can also use Amazon CloudFront with origin servers running on Amazon EC2 instances or with any other custom origin.

Reference: <http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/distribution-web.html>

NEW QUESTION 432

You have been asked to set up monitoring of your network and you have decided that Cloudwatch would be the best service to use. Amazon CloudWatch monitors

your Amazon Web Services (AWS) resources and the applications you run on AWS in real-time. You can use CloudWatch to collect and track metrics, which are the variables you want to measure for your resources and applications. Which of the following items listed can AWS Cloudwatch monitor?

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

Answer: B

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

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

NEW QUESTION 436

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 437

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 438

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 440

You have a Business support plan with AWS. One of your EC2 instances is running Microsoft Windows Server 2008 R2 and you are having problems with the software. Can you receive support from AWS for this software?

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

Answer: A

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

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

NEW QUESTION 445

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

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

Answer: A

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

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

CloudWatch Metrics

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

Ganglia

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

Reference:

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

NEW QUESTION 450

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

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

Answer: A

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

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

NEW QUESTION 455

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 457

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 460

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

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

Answer: B

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

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

NEW QUESTION 461

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 464

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

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

Answer: C

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

Benefits

Enhanced Durability

Multi-AZ deployments for the MySQL, Oracle, and PostgreSQL engines utilize synchronous physical replication to keep data on the standby up-to-date with the primary. Multi-AZ deployments for the SQL Server engine use synchronous logical replication to achieve the same result, employing SQL

Server-native Mrroring technology. Both approaches safeguard your data in the event of a DB Instance failure or loss of an Availability Zone.

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

Amazon Aurora employs a highly durable, SSD-backed virtualized storage layer purpose-built for

database workloads. Amazon Aurora automatically replicates your volume six ways, across three Availability Zones. Amazon Aurora storage is fault-tolerant, transparently handling the loss of up to two copies of data without affecting database write availability and up to three copies without affecting read availability.

Amazon Aurora storage is also self-healing. Data blocks and disks are continuously scanned for errors and replaced automatically.

Increased Availability

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

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

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

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

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

No Administrative Intervention

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

Failover conditions

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

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

Storage failure on primary

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

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

NEW QUESTION 469

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

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

Answer: A

Explanation: Point-In-Time Recovery

In addition to the daily automated backup, Amazon RDS archives database change logs. This enables you to recover your database to any point in time during the backup retention period, up to the last five minutes of database usage.

Amazon RDS stores multiple copies of your data, but for Single-AZ DB instances these copies are stored in a single availability zone. If for any reason a Single-AZ DB instance becomes unusable, you can use point-in-time recovery to launch a new DB instance with the latest restorable data. For more information on working with point-in-time recovery, go to Restoring a DB Instance to a Specified Time.

Note

Multi-AZ deployments store copies of your data in different Availability Zones for greater levels of data durability. For more information on Multi-AZ deployments, see High Availability (Multi-AZ).

NEW QUESTION 471

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 S3 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 473

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 477

Company B is launching a new game app for mobile devices. Users will log into the game using their existing social media account to streamline data capture. Company B would like to directly save player data and scoring information from the mobile app to a DynamoDB table named Score Data. When a user saves their game the progress data will be stored to the Game state 53 bucket. What is the best approach for storing data to DynamoDB and S3?

- A. Use an EC2 Instance that is launched with an EC2 role providing access to the Score Data DynamoDB table and the Game State S3 bucket that communicates with the mobile app via web services.
- B. Use temporary security credentials that assume a role providing access to the Score Data DynamoDB table and the Game State S3 bucket using web identity federation.
- C. Use Login with Amazon allowing users to sign in with an Amazon account providing the mobile app with access to the Score Data DynamoDB table and the Game State S3 bucket.
- D. Use an IAM user with access credentials assigned a role providing access to the Score Data DynamoDB table and the Game State S3 bucket for distribution with the mobile app.

Answer: B

Explanation: Web Identity Federation

Imagine that you are creating a mobile app that accesses AWS resources, such as a game that runs on a mobile device and stores player and score information using Amazon S3 and DynamoDB. When you write such an app, you'll make requests to AWS services that must be signed with an AWS access key. However, we strongly recommend that you do not embed or distribute long-term AWS credentials with apps that a user downloads to a device, even in an encrypted store. Instead, build your app so that it requests temporary AWS security credentials dynamically when needed using web identity federation. The supplied temporary credentials map to an AWS role that has only the permissions needed to perform the tasks required by the mobile app.

With web identity federation, you don't need to create custom sign-in code or manage your own user identities. Instead, users of your app can sign in using a well-known identity provider (IdP) - such as Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible IdP, receive an authentication token, and then exchange that token for temporary security credentials in AWS that map to an IAM role with permissions to use the resources in your AWS account. Using an IdP helps you keep your AWS account secure, because you don't have to embed and distribute long-term security credentials with your application.

For most scenarios, we recommend that you use Amazon Cognito because it acts as an identity broker and does much of the federation work for you. For details, see the following section, Using Amazon Cognito for Mobile Apps.

If you don't use Amazon Cognito, then you must write code that interacts with a web IdP (Login with Amazon, Facebook, Google, or any other OIDC-compatible IdP) and then calls the Assume Role With Web Identity API to trade the authentication token you get from those IdPs for AWS temporary security credentials. If you have already used this approach for existing apps, you can continue to use it.

Using Amazon Cognito for Mobile Apps

The preferred way to use web identity federation is to use Amazon Cognito. For example, Adele the developer is building a game for a mobile device where user data such as scores and profiles is stored in Amazon S3 and Amazon DynamoDB. Adele could also store this data locally on the device and use Amazon Cognito to keep it synchronized across devices. She knows that for security and maintenance reasons, long-term AWS security credentials should not be distributed with the game. She also knows that the game might have a large number of users. For all of these reasons, she does not want to create new user identities in IAM for each player. Instead, she builds the game so that users can sign in using an identity that they've already established with a well-known identity provider, such as Login with Amazon, Facebook, Google, or any OpenID Connect (OIDC)-compatible identity provider.

Her game can take advantage of the authentication mechanism from one of these providers to validate the user's identity.


To enable the mobile app to access her AWS resources, Adele first registers for a developer ID with her chosen IdPs. She also configures the application with each of these providers. In her AWS account that contains the Amazon S3 bucket and DynamoDB table for the game, Adele uses Amazon Cognito to create IAM roles that precisely define permissions that the game needs. If she is using an OIDC IdP, she also creates an IAM OIDC identity provider entity to establish trust between her AWS account and the IdP.

In the app's code, Adele calls the sign-in interface for the IdP that she configured previously. The IdP handles all the details of letting the user sign in, and the app gets an OAuth access token or OIDC ID token from the provider. Adele's app can trade this authentication information for a set of temporary security credentials that consist of an AWS access key ID, a secret access key, and a session token.

The app can then use these credentials to access web services offered by AWS. The app is limited to the permissions that are defined in the role that it assumes. The following figure shows a simplified flow for how this might work, using Login with Amazon as the IdP.

For Step 2, the app can also use Facebook, Google, or any OIDC-compatible identity provider, but that's not shown here.

Sample workflow using Amazon Cognito to federate users for a mobile application



A customer starts your app on a mobile device. The app asks the user to sign in. The app uses Login with Amazon resources to accept the user's credentials. The app uses Cognito APIs to exchange the Login with Amazon ID token for a Cognito token. The app requests temporary security credentials from AWS STS, passing the Cognito token.

The temporary security credentials can be used by the app to access any AWS resources required by the app to operate. The role associated with the temporary security credentials and its assigned policies determines what can be accessed.

Use the following process to configure your app to use Amazon Cognito to authenticate users and give your app access to AWS resources. For specific steps to accomplish this scenario, consult the documentation for Amazon Cognito.

(Optional) Sign up as a developer with Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible identity provider and configure

one or more apps with the provider. This step is optional because Amazon Cognito also supports unauthenticated (guest) access for your users. Go to Amazon Cognito in the AWS IAM Management Console. Use the Amazon Cognito wizard to create an identity pool, which is a container that Amazon Cognito uses to keep end user identities organized for your apps. You can share identity pools between apps. When you set up an identity pool, Amazon Cognito creates one or two IAM roles (one for authenticated identities, and one for unauthenticated "guest" identities) that define permissions for Amazon Cognito users. Download and integrate the AWS SDK for iOS or the AWS SDK for Android with your app, and import the files required to use Amazon Cognito. Create an instance of the Amazon Cognito credentials provider, passing the identity pool ID, your AWS account number, and the Amazon Resource Name (ARN) of the roles that you associated with the identity pool. The Amazon Cognito wizard in the AWS Management Console provides sample code to help you get started. When your app accesses an AWS resource, pass the credentials provider instance to the client object, which passes temporary security credentials to the client. The permissions for the credentials are based on the role or roles that you defined earlier.

NEW QUESTION 480

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

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

Answer: B

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

What can I do with Amazon SQS?

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

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

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

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

NEW QUESTION 481

You need a persistent and durable storage to trace call activity of an IVR (Interactive Voice Response) system. Call duration is mostly in the 2-3 minutes timeframe. Each traced call can be either active or terminated. An external application needs to know each minute the list of currently active calls, which are usually a few calls/second. Put once per month there is a periodic peak up to 1000 calls/second for a few hours. The system is open 24/7 and any downtime should be avoided.

Historical data is periodically archived to files. Cost saving is a priority for this project.

What database implementation would better fit this scenario, keeping costs as low as possible?

- A. Use RDS Multi-AZ with two tables, one for "Active calls" and one for "Terminated calls". In this way the "Active calls" table is always small and effective to access.
- B. Use DynamoDB with a "Calls" table and a Global Secondary Index on a "IsActive" attribute that is present for active calls only. In this way the Global Secondary index is sparse and more effective.
- C. Use DynamoDB with a "Calls" table and a Global secondary index on a "State" attribute that can equal to "active" or "terminated" in this way the Global Secondary index can be used for all items in the table.
- D. Use RDS Multi-AZ with a "CALLS" table and an Indexed "STATE" field that can be equal to "ACTIVE" or "TERMINATED". In this way the SQL query is optimized by the use of the Index.

Answer: A

NEW QUESTION 483

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

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

Answer: AC

Explanation: Reference:

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

NEW QUESTION 484

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 489

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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