

AWS-Solution-Architect-Associate Dumps

Amazon AWS Certified Solutions Architect - Associate

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



NEW QUESTION 1

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

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

Answer: A

Explanation:

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

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-ug.pdf>

NEW QUESTION 2

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

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

Answer: D

Explanation:

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

AWS recommends Amazon Redshift for customers who have a combination of needs, such as: High performance at scale as data and query complexity grows
Desire to prevent reporting and analytic processing from interfering with the performance of OLTP workloads

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

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

NEW QUESTION 3

A 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 4

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 5

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

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

Answer: D

Explanation:

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

NEW QUESTION 6

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 7

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

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

Answer: B

Explanation:

By default, all AWS accounts are limited to 5 Elastic IP addresses per region for each AWS account, because public (IPv4) Internet addresses are a scarce public resource. AWS strongly encourages you to use an EIP primarily for load balancing use cases, and use DNS hostnames for all other inter-node communication. If you feel your architecture warrants additional EIPs, you would need to complete the Amazon EC2 Elastic IP Address Request Form and give reasons as to your need for additional addresses. Reference:

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

NEW QUESTION 8

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 9

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

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

Answer: C

Explanation:

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

Columnar storage for database tables is an important factor in optimizing analytic query performance because it drastically reduces the overall disk I/O

requirements and reduces the amount of data you need to load from disk.

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

NEW QUESTION 10

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 10

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

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

Answer: B

Explanation:

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

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

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

NEW QUESTION 14

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 15

In DynamoDB, could you use IAM to grant access to Amazon DynamoDB resources and API actions?

- A. In DynamoDB there is no need to grant access
- B. Depended to the type of access
- C. No
- D. Yes

Answer: D

Explanation:

Amazon DynamoDB integrates with AWS Identity and Access Management (IAM). You can use AWS IAM to grant access to Amazon DynamoDB resources and API actions. To do this, you first write an AWS IAM policy, which is a document that explicitly lists the permissions you want to grant. You then attach that policy to an AWS IAM user or role.

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

NEW QUESTION 18

As AWS grows, most of your clients' main concerns seem to be about security, especially when all of their competitors also seem to be using AWS. One of your clients asks you whether having a competitor who hosts their EC2 instances on the same physical host would make it easier for the competitor to hack into the client's data. Which of the following statements would be the best choice to put your client's mind at rest?

- A. Different instances running on the same physical machine are isolated from each other via a 256-bit Advanced Encryption Standard (AES-256).
- B. Different instances running on the same physical machine are isolated from each other via the Xen hypervisor and via a 256-bit Advanced Encryption Standard (AES-256).
- C. Different instances running on the same physical machine are isolated from each other via the Xen hypervisor.
- D. Different instances running on the same physical machine are isolated from each other via IAM permissions.

Answer: C

Explanation:

Amazon Elastic Compute Cloud (EC2) is a key component in Amazon's Infrastructure as a Service (IaaS), providing resizable computing capacity using server instances in AWS's data centers. Amazon EC2 is designed to make web-scale computing easier by enabling you to obtain and configure capacity with minimal friction.

You create and launch instances, which are collections of platform hardware and software. Different instances running on the same physical machine are isolated from each other via the Xen hypervisor.

Amazon is active in the Xen community, which provides awareness of the latest developments. In addition, the AWS firewall resides within the hypervisor layer, between the physical network interface and the instance's virtual interface. All packets must pass through this layer, thus an instance's neighbors have no more access to that instance than any other host on the Internet and can be treated as if they are on separate physical hosts. The physical RAM is separated using similar mechanisms.

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

NEW QUESTION 22

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 24

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

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

Answer: B

Explanation:

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

NEW QUESTION 26

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 31

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 33

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 35

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 36

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 38

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 41

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 42

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 44

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

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

Answer: B

Explanation:

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

AWS charges less where their costs are less.

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

NEW QUESTION 48

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 52

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 57

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 58

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 62

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

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

Answer: C

Explanation:

IAM users need passwords in order to access the AWS Management Console. (They do not need passwords if they will access AWS resources programmatically by using the CLI, AWS SDKs, or the APIs.)

You can use a password policy to do these things: Set a minimum password length.

Require specific character types, including uppercase letters, lowercase letters, numbers, and non-alphanumeric characters. Be sure to remind your users that passwords are case sensitive. Allow all IAM users to change their own passwords.

Require IAM users to change their password after a specified period of time (enable password expiration). Prevent IAM users from reusing previous passwords.

Force IAM users to contact an account administrator when the user has allowed his or her password to expire.

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

NEW QUESTION 64

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 67

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 71

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 73

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 78

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 81

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 86

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 91

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 traffi

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 95

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 100

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 102

A government client needs you to set up secure cryptographic key storage for some of their extremely confidential data. You decide that the AWS CloudHSM is the best service for this. However, there seem to be a few pre-requisites before this can happen, one of those being a security group that has certain ports open. Which of the following is correct in regards to those security groups?

- A. A security group that has port 22 (for SSH) or port 3389 (for RDP) open to your network.
- B. A security group that has no ports open to your network.
- C. A security group that has only port 3389 (for RDP) open to your network.
- D. A security group that has only port 22 (for SSH) open to your network.

Answer: A

Explanation:

AWS CloudHSM provides secure cryptographic key storage to customers by making hardware security modules (HSMs) available in the AWS cloud.

AWS CloudHSM requires the following environment before an HSM appliance can be provisioned. A virtual private cloud (VPC) in the region where you want the AWS CloudHSM service.

One private subnet (a subnet with no Internet gateway) in the VPC. The HSM appliance is provisioned into this subnet.

One public subnet (a subnet with an Internet gateway attached). The control instances are attached to this subnet.

An AWS Identity and Access Management (IAM) role that delegates access to your AWS resources to AWS CloudHSM.

An EC2 instance, in the same VPC as the HSM appliance, that has the SafeNet client software installed. This instance is referred to as the control instance and is used to connect to and manage the HSM appliance.

A security group that has port 22 (for SSH) or port 3389 (for RDP) open to your network. This security group is attached to your control instances so you can access them remotely.

NEW QUESTION 104

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 107

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 HSM.
- 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 111

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 115

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 120

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.aws.amazon.com/AWSEC2/latest/UserGuide/using-instance-addressing.html>

NEW QUESTION 124

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

Inbound SSH traffic.

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

- A. Ports 22,1433,3389

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 127

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 130

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 133

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

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

Answer: B

Explanation:

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

NEW QUESTION 138

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

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

Answer: B

Explanation:

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

Reference:

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

NEW QUESTION 143

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 148

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 150

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 155

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 157

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 161

You have launched an EC2 instance with four (4) 500GB EBS Provisioned IOPS volumes attached The EC2 Instance Is EBS-Optimized and supports 500 Mbps throughput between EC2 and EBS The two EBS volumes are configured as a single RAID 0 device, and each Provisioned IOPS volume is provisioned with 4.000 IOPS (4 000 16KB reads or writes) for a total of 16.000 random IOPS on the instance The EC2 Instance initially delivers the expected 16 000 IOPS random read and write performance Sometime later in order to increase the total random I/O performance of the instance, you add an additional two 500 GB EBS Provisioned IOPS volumes to the RAID Each volume Is provisioned to 4.000 IOPs like the original four for a total of 24.000 IOPS on the EC2 instance Monitoring shows that the EC2 instance CPU utilization increased from 50% to 70%. but the total random IOPS measured at the instance level does not increase at all. What is the problem and a valid solution?

- A. Larger storage volumes support higher Provisioned IOPS rates: increase the provisioned volumestorage of each of the 6 EBS volumes to ITB
- B. The EBS-Optimized throughput limits the total IOPS that can be utilized use an EBS-Optimized instance that provides larger throughput.
- C. Small block sizes cause performance degradation, limiting the I/O throughput, configure the instance device driver and file system to use 64KB blocks to increase throughput.
- D. RAID 0 only scales linearly to about 4 devices, use RAID 0 with 4 EBS Provisioned IOPS volumes but increase each Provisioned IOPS EBS volume to 6.000 IOPS.
- E. The standard EBS instance root volume limits the total IOPS rate, change the instant root volume to also be a 500GB 4.000 Provisioned IOPS volume.

Answer: E

NEW QUESTION 165

You need a persistent and durable storage to trace call actMty 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 166

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

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

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

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

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

Answer: C

NEW QUESTION 168

Your company previously configured a heavily used, dynamically routed VPN connection between your on-premises data center and AWS. You recently provisioned a DirectConnect connection and would like to start using the new connection. After configuring DirectConnect settings in the AWS Console, which of the following options win provide the most seamless transition for your users?

- A. Delete your existing VPN connection to avoid routing loops configure your DirectConnect router with the appropriate settings and verify network traffic is leveraging DirectConnect.
- B. Configure your DirectConnect router with a higher BGP priority than your VPN router, verify network traffic is leveraging Directconnect and then delete your existing VPN connection.
- C. Update your VPC route tables to point to the DirectConnect connection configure your DirectConnect router with the appropriate settings verify network traffic is leveraging DirectConnect and then delete the VPN connection.
- D. Configure your DirectConnect router, update your VPC route tables to point to the DirectConnect connection, configure your VPN connection with a higher BGP point
- E. And verify network traffic is leveraging the DirectConnect connection.

Answer: D

NEW QUESTION 171

You are designing the network infrastructure for an application server in Amazon VPC Users will access all the application instances from the Internet as well as from an on-premises network The on-premises network is connected to your VPC over an AWS Direct Connect link. How would you design routing to meet the above requirements?

- A. Configure a single routing Table with a default route via the Internet gateway Propagate a default route via BGP on the AWS Direct Connect customer route
- B. Associate the routing table with all VPC subnets.
- C. Configure a single routing table with a default route via the internet gateway Propagate specific routes for the on-premises networks via BGP on the AWS Direct Connect customer router Associate the routing table with all VPC subnets.
- D. Configure a single routing table with two default routes: one to the internet via an Internet gateway the other to the on-premises network via the VPN gateway use this routing table across all subnets in your VPC,
- E. Configure two routing tables one that has a default route via the Internet gateway and another that has a default route via the VPN gateway Associate both routing tables with each VPC subnet.

Answer: A

NEW QUESTION 172

You are implementing AWS Direct Connect. You intend to use AWS public service end points such as Amazon S3, across the AWS Direct Connect link. You want other Internet traffic to use your existing link to an Internet Service Provider. What is the correct way to configure AWS Direct connect for access to services such as Amazon S3?

- A. Configure a public Interface on your AWS Direct Connect link Configure a static route via your AWS Direct Connect link that points to Amazon S3 Advertise a default route to AWS using BGP.
- B. Create a private interface on your AWS Direct Connect link
- C. Configure a static route via your AWS Direct connect link that points to Amazon S3 Configure specific routes to your network in your VPC,
- D. Create a public interface on your AWS Direct Connect link Redistribute BGP routes into your existing routing infrastructure advertise specific routes for your network to AWS.
- E. Create a private interface on your AWS Direct connect link
- F. Redistribute BGP routes into your existing routing infrastructure and advertise a default route to AWS.

Answer: C

NEW QUESTION 174

You have deployed a three-tier web application in a VPC with a CIDR block of 10.0.0.0/28 You initially deploy two web servers, two application servers, two database servers and one NAT instance for a total of seven EC2 instances The web, Application and database servers are deployed across two availability zones (AZs). You also deploy an ELB in front of the two web servers, and use Route53 for DNS Web traffic gradually increases in the first few days following the deployment, so you attempt to double the number of instances in each tier of the application to handle the new load unfortunately some of these new instances fail to launch.

Which of the following could be the root cause? (Choose 2 answers)

- A. AWS reserves the first and the last private IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances
- B. The Internet Gateway (IGW) of your VPC has scaled-up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches
- C. The ELB has scaled-up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches
- D. AWS reserves one IP address in each subnet's CIDR block for Route53 so you do not have enough addresses left to launch all of the new EC2 instances
- E. AWS reserves the first four and the last IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances

Answer: CE

NEW QUESTION 176

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

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

Answer: ADF

Explanation:

AWS Direct Connect

AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS, your datacenter, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet based connections.

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

What is AWS Direct Connect?

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

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

Requirements

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

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

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

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

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

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

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

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

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

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

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

The public routes that you will advertise over BGP.

NEW QUESTION 180

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

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

Answer: C

Explanation:

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

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

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

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

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

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

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

Database Layer: A few options:

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

Use RDS to run MySQL/ Postgres

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

App Layer:

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

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

Preventing Abuse: When you put something in a public S3 bucket, people will hot-link it from their web pages. If you want to prevent that, your app on the EC2 box

can generate signed links to S3 that expire in a few hours. Now everyone will be forced to go thru the app, and the app can apply rate limiting, etc. Saving money: If you don't mind having downtime: run everything in one AZ (both DBs and EC2s). You can always add servers and AZs down the road, as long as it's architected to be stateless. In fact, you should use multiple regions if you want it to be really robust. use Reduced Redundancy in S3 to save a few hundred bucks per month (Someone will have to "go fix it" every time it breaks, including having an off-line copy to repair S3.) Buy Reserved Instances on your EC2 boxes to make them cheaper. (Start with the RI market and buy a partially used one to get started.) It's just a coupon saying "if you run this type of box in this AZ, you will save on the per-hour costs." You can get 1/2 to 1/3 off easily. Rewrite the application to use less memory and CPU -that way you can run on fewer/ smaller boxes. (May or may not be worth the investment.) If your app will be used very infrequently, you will save a lot of money by using Lambda. I'd be worried that it would be quite slow if you tried to run a Java application on it though .. We're missing some information like load, latency expectations from search, indexing speed, size of the search index, etc. But with what you've given us, I would go with S3 as the storage for the files (S3 rocks. It is really, really awesome). If you're stuck with the commercial search application, then on EC2 instances with autoscaling and an ELB. If you are allowed an alternative search engine, Elasticsearch is probably your best bet. I'd run it on EC2 instead of the AWS Elasticsearch service, as IMHO it's not ready yet. Don't autoscale Elasticsearch automatically though, it'll cause all sorts of issues. I have zero experience with CloudSearch so I can't comment on that. Regardless of which option, I'd use CloudFormation for all of it.

NEW QUESTION 185

You're running an application on-premises due to its dependency on non-x86 hardware and want to use AWS for data backup. Your backup application is only able to write to POSIX-compatible block-based storage. You have 140TB of data and would like to mount it as a single folder on your file server. Users must be able to access portions of this data while the backups are taking place. What backup solution would be most appropriate for this use case?

- A. Use Storage Gateway and configure it to use Gateway-Cached volumes.
- B. Configure your backup software to use S3 as the target for your data backups.
- C. Configure your backup software to use Glacier as the target for your data backups.
- D. Use Storage Gateway and configure it to use Gateway-Stored volume

Answer: A

Explanation:**Gateway-Cached Volume Architecture**

Gateway-cached volumes let you use Amazon Simple Storage Service (Amazon S3) as your primary data storage while retaining frequently accessed data locally in your storage gateway. Gateway-cached volumes minimize the need to scale your on-premises storage infrastructure, while still providing your applications with low-latency access to their frequently accessed data. You can create storage volumes up to 32 TiB in size and attach to them as iSCSI devices from your on-premises application servers. Your gateway stores data that you write to these volumes in Amazon S3 and retains recently read data in your on-premises storage gateway's cache and upload buffer storage.

Gateway-cached volumes can range from 1 GiB to 32 TiB in size and must be rounded to the nearest GiB. Each gateway configured for gateway-cached volumes can support up to 32 volumes for a total maximum storage volume of 1,024 TiB (1 PiB).

In the gateway-cached volume solution, AWS Storage Gateway stores all your on-premises application data in a storage volume in Amazon S3.

The following diagram provides an overview of the AWS Storage Gateway-cached volume deployment.

After you've installed the AWS Storage Gateway software appliance-the virtual machine (VM)-on a host in your data center and activated it, you can use the AWS Management Console to provision storage

volumes backed by Amazon S3. You can also provision storage volumes programmatically using the AWS Storage Gateway API or the AWS SDK libraries. You then mount these storage volumes to your on-premises application servers as iSCSI devices.

You also allocate disks on-premises for the VM. These on-premises disks serve the following purposes: Disks for use by the gateway as cache storage - As your applications write data to the storage volumes in AWS, the gateway initially stores the data on the on-premises disks referred to as cache storage before uploading the data to Amazon S3. The cache storage acts as the on-premises durable store for data that is waiting to upload to Amazon S3 from the upload buffer.

The cache storage also lets the gateway store your application's recently accessed data on-premises for low-latency access. If your application requests data, the gateway first checks the cache storage for the data before checking Amazon S3.

You can use the following guidelines to determine the amount of disk space to allocate for cache storage. Generally, you should allocate at least 20 percent of your existing file store size as cache storage. Cache storage should also be larger than the upload buffer. This latter guideline helps ensure cache storage is large enough to persistently hold all data in the upload buffer that has not yet been uploaded to Amazon S3.

Disks for use by the gateway as the upload buffer - To prepare for upload to Amazon S3, your gateway also stores incoming data in a staging area, referred to as an upload buffer. Your gateway uploads this buffer data over an encrypted Secure Sockets Layer (SSL) connection to AWS, where it is stored encrypted in Amazon S3.

You can take incremental backups, called snapshots, of your storage volumes in Amazon S3. These point-in-time snapshots are also stored in Amazon S3 as Amazon EBS snapshots. When you take a new snapshot, only the data that has changed since your last snapshot is stored. You can initiate snapshots on a scheduled or one-time basis. When you delete a snapshot, only the data not needed for any other snapshots is removed.

You can restore an Amazon EBS snapshot to a gateway storage volume if you need to recover a backup of your data. Alternatively, for snapshots up to 16 TiB in size, you can use the snapshot as a starting point for a new Amazon EBS volume. You can then attach this new Amazon EBS volume to an Amazon EC2 instance.

All gateway-cached volume data and snapshot data is stored in Amazon S3 encrypted at rest using server-side encryption (SSE). However, you cannot access this data with the Amazon S3 API or other tools such as the Amazon S3 console.

NEW QUESTION 188

You are the new IT architect in a company that operates a mobile sleep tracking application

When activated at night, the mobile app is sending collected data points of 1 kilobyte every 5 minutes to your backend

The backend takes care of authenticating the user and writing the data points into an Amazon DynamoDB table.

Every morning, you scan the table to extract and aggregate last night's data on a per user basis, and store the results in Amazon S3.

Users are notified via Amazon SNS mobile push notifications that new data is available, which is parsed and visualized by the mobile app. Currently you have around 100k users who are mostly based out of North America.

You have been tasked to optimize the architecture of the backend system to lower cost. What would you recommend? (Choose 2 answers)

- A. Create a new Amazon DynamoDB table each day and drop the one for the previous day after its data is on Amazon S3.
- B. Have the mobile app access Amazon DynamoDB directly instead of JSON files stored on Amazon S3.
- C. Introduce an Amazon SQS queue to buffer writes to the Amazon DynamoDB table and reduce provisioned write throughput.
- D. Introduce Amazon ElastiCache to cache reads from the Amazon DynamoDB table and reduce provisioned read throughput.
- E. Write data directly into an Amazon Redshift cluster replacing both Amazon DynamoDB and Amazon S3.

Answer: BD

NEW QUESTION 189

Your website is serving on-demand training videos to your workforce. Videos are uploaded monthly in high resolution MP4 format. Your workforce is distributed globally often on the move and using company-provided tablets that require the HTTP Live Streaming (HLS) protocol to watch a video. Your company has no video transcoding expertise and it required you may need to pay for a consultant.

How do you implement the most cost-efficient architecture without compromising high availability and quality of video delivery'?

- A. A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queue
- B. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few day
- C. CloudFront to serve HLS transcoded videos from EC2.
- D. Elastic Transcoder to transcode original high-resolution MP4 videos to HL
- E. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few day
- F. CloudFront to serve HLS transcoded videos from EC2.
- G. Elastic Transcoder to transcode original high-resolution NIP4 videos to HL
- H. 53 to host videos with Lifecycle Management to archive original files to Glacier after a few day
- I. CloudFront to serve HLS transcoded videos from 53.
- J. A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queue
- K. 53 to host videos with Lifecycle Management to archive all files to Glacier after a few day
- L. CloudFront to serve HLS transcoded videos from Glacier.

Answer: C

NEW QUESTION 193

You currently operate a web application in the AWS US-East region. The application runs on an autoscaled layer of EC2 instances and an RDS Multi-AZ database. Your IT security compliance officer has tasked you to develop a reliable and durable logging solution to track changes made to your EC2, IAM, and RDS resources. The solution must ensure the integrity and confidentiality of your log data. Which of these solutions would you recommend?

- A. Create a new CloudTrail trail with one new S3 bucket to store the logs and with the global services option selected. Use IAM roles, S3 bucket policies, and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs.
- B. Create a new CloudTrail trail with one new S3 bucket to store the logs. Configure SNS to send log file delivery notifications to your management system. Use IAM roles and S3 bucket policies on the S3 bucket that stores your logs.
- C. Create a new CloudTrail trail with an existing S3 bucket to store the logs and with the global services option selected. Use S3 ACLs and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs.
- D. Create three new CloudTrail trails with three new S3 buckets to store the logs: one for the AWS Management console, one for AWS SDKs, and one for command line tools. Use IAM roles and S3 bucket policies on the S3 buckets that store your logs.

Answer: A

NEW QUESTION 196

You have an application running on an EC2 instance which will allow users to download files from a private S3 bucket using a pre-assigned URL. Before generating the URL, the application should verify the existence of the file in S3.

How should the application use AWS credentials to access the S3 bucket securely?

- A. Use the AWS account access keys; the application retrieves the credentials from the source code of the application.
- B. Create an IAM user for the application with permissions that allow list access to the S3 bucket. Launch the instance as the IAM user and retrieve the IAM user's credentials from the EC2 instance user data.
- C. Create an IAM role for EC2 that allows list access to objects in the S3 bucket.
- D. Launch the instance with the role, and retrieve the role's credentials from the EC2 instance metadata.
- E. Create an IAM user for the application with permissions that allow list access to the S3 bucket.
- F. The application retrieves the IAM user credentials from a temporary directory with permissions that allow read access only to the application user.

Answer: C

NEW QUESTION 197

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

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

Answer: B

Explanation:

Reference: <http://www.slideshare.net/AmazonWebServices/aws-webcast-introduction-to-amazon-kinesis>

NEW QUESTION 198

You are implementing a URL whitelisting system for a company that wants to restrict outbound HTTP/S connections to specific domains from their EC2-hosted applications. You deploy a single EC2 instance running proxy software and configure it to accept traffic from all subnets and EC2 instances in the VPC. You configure the proxy to only pass through traffic to domains that you define in its whitelist configuration. You have a nightly maintenance window of 10 minutes where all instances fetch new software updates. Each update is about 200MB in size and there are 500 instances in the VPC that routinely fetch updates. After a few days you notice that some machines are failing to successfully download some, but not all of their updates within the maintenance window. The download URLs used for these updates are correctly listed in the proxy's whitelist configuration and you are able to access them manually using a web browser on the instances. What might be happening? (Choose 2 answers)

- A. You are running the proxy on an undersized EC2 instance type so network throughput is not sufficient for all instances to download their updates in time.
- B. You are running the proxy on a sufficiently-sized EC2 instance in a private subnet and its network throughput is being throttled by a NAT running on an

undersized EC2 instance.

- C. The route table for the subnets containing the affected EC2 instances is not configured to direct network traffic for the software update locations to the proxy.
- D. You have not allocated enough storage to the EC2 instance running the proxy so the network buffer is filling up, causing some requests to fail.
- E. You are running the proxy in a public subnet but have not allocated enough EIPs to support the needed network throughput through the Internet Gateway (IGW).

Answer: AB

NEW QUESTION 202

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

- A. Use Cloud Front and an Elastic Load balancer in front of an auto-scaled set of web servers, the web servers will first call the Login With Amazon service to authenticate the user then process the users vote and store the result into a multi-AZ Relational Database Service instance.
- B. Use CloudFront and the static website hosting feature of S3 with the Javascript SDK to call the Login With Amazon service to authenticate the user, use IAM Roles to gain permissions to a DynamoDB table to store the users vote.
- C. Use Cloud Front and an Elastic Load Balancer in front of an auto-scaled set of web servers, the web servers will first call the Login with Amazon service to authenticate the user, the web servers will process the users vote and store the result into a DynamoDB table using IAM Roles for EC2 instances to gain permissions to the DynamoDB table.
- D. Use Cloud Front and an Elastic Load Balancer in front of an auto-scaled set of web servers, the web servers will first call the Login With Amazon service to authenticate the user, the web servers will process the users vote and store the result into an SQS queue using IAM Roles for EC2 Instances to gain permissions to the SQS queue.
- E. With Amazon service to authenticate the user, the web servers will process the users vote and store the result into an SQS queue using IAM Roles for EC2 Instances to gain permissions to the SQS queue.
- F. A set of application servers will then retrieve the items from the queue and store the result into a DynamoDB table.

Answer: D

NEW QUESTION 203

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 208

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

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

Answer: C

NEW QUESTION 210

You must assign each instance to at least one security group

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

Answer: A

NEW QUESTION 214

All Amazon EC2 instances are assigned two IP addresses at launch, out of which one can only be reached from within the Amazon EC2 network?

- A. Multiple IP address
- B. Public IP address
- C. Private IP address
- D. Elastic I P Address

Answer: C

NEW QUESTION 218

What is the Reduced Redundancy option in Amazon S3?

- A. Less redundancy for a lower cost.
- B. It doesn't exist in Amazon S3, but in Amazon EBS.
- C. It allows you to destroy any copy of your files outside a specific jurisdiction.
- D. It doesn't exist at all

Answer: A

NEW QUESTION 221

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

- A. Delete On Termination
- B. Remove On Deletion
- C. Remove On Termination
- D. Terminate On Deletion

Answer: A

NEW QUESTION 222

What does Amazon Elastic Beanstalk provide?

- A. A scalable storage appliance on top of Amazon Web Services.
- B. An application container on top of Amazon Web Services.
- C. A service by this name doesn't exist.
- D. A scalable cluster of EC2 instances

Answer: B

NEW QUESTION 224

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

- A. TRUE
- B. FALSE

Answer: A

NEW QUESTION 228

Amazon RDS DB snapshots and automated backups are stored in

- A. Amazon S3
- B. Amazon ECS Volume
- C. Amazon RDS
- D. Amazon EMR

Answer: A

NEW QUESTION 231

Using Amazon CloudWatch's Free Tier, what is the frequency of metric updates which you receive?

- A. 5 minutes
- B. 500 milliseconds.
- C. 30 seconds
- D. 1 minute

Answer: A

NEW QUESTION 232

What does the following command do with respect to the Amazon EC2 security groups? `ec2-revoke RevokeSecurityGroup Ingress`

- A. Removes one or more security groups from a rule.
- B. Removes one or more security groups from an Amazon EC2 instance.
- C. Removes one or more rules from a security group.
- D. Removes a security group from our account

Answer: C

NEW QUESTION 237

Can a 'user' be associated with multiple AWS accounts?

- A. No
- B. Yes

Answer: A

NEW QUESTION 242

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

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

Answer: D

NEW QUESTION 243

IAM provides several policy templates you can use to automatically assign permissions to the groups you create. The _ policy template gives the Admins group permission to access all account resources, except your AWS account information

- A. Read Only Access
- B. Power User Access
- C. AWS Cloud Formation Read Only Access
- D. Administrator Access

Answer: D

NEW QUESTION 244

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

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

Answer: B

NEW QUESTION 248

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

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

Answer: D

NEW QUESTION 250

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

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

Answer: A

NEW QUESTION 255

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

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

Answer: D

NEW QUESTION 260

In the 'Detailed' monitoring data available for your Amazon EBS volumes, Provisioned IOPS volumes automatically send _ minute metrics to Amazon CloudWatch.

- A. 3
- B. 1
- C. 5
- D. 2

Answer: B

NEW QUESTION 262

Are Reserved Instances available for Multi-AZ Deployments?

- A. Only for Cluster Compute instances
- B. Yes for all instance types
- C. Only for M3 instance types
- D. No

Answer: B

NEW QUESTION 263

If I modify a DB Instance or the DB parameter group associated with the instance, should I reboot the instance for the changes to take effect?

- A. No
- B. Yes

Answer: B

NEW QUESTION 266

Making your snapshot public shares all snapshot data with everyone. Can the snapshots with AWS Marketplace product codes be made public?

- A. No
- B. Yes

Answer: B

NEW QUESTION 271

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

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

Answer: A

NEW QUESTION 273

What does Amazon CloudFormation provide?

- A. The ability to setup Autoscaling for Amazon EC2 instances.
- B. None of these.
- C. A templated resource creation for Amazon Web Services.
- D. A template to map network resources for Amazon Web Service

Answer: D

NEW QUESTION 275

What are the four levels of AWS Premium Support?

- A. Basic, Developer, Business, Enterprise
- B. Basic, Startup, Business, Enterprise
- C. Free, Bronze, Silver, Gold
- D. All support is free

Answer: A

NEW QUESTION 277

What can I access by visiting the URL: <http://status.aws.amazon.com/>?

- A. Amazon Cloud Watch
- B. Status of the Amazon RDS DB
- C. AWS Service Health Dashboard
- D. AWS Cloud Monitor

Answer: C

NEW QUESTION 280

Please select the Amazon EC2 resource which cannot be tagged.

- A. images (AMIs, kernels, RAM disks)
- B. Amazon EBS volumes
- C. Elastic IP addresses
- D. VPCs

Answer: C

NEW QUESTION 285

The new DB Instance that is created when you promote a Read Replica retains the backup window period.

- A. TRUE
- B. FALSE

Answer: A

NEW QUESTION 288

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

- A. FALSE
- B. TRUE

Answer: A

NEW QUESTION 291

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

- A. Only if instructed to when created
- B. Yes
- C. No

Answer: B

NEW QUESTION 293

After an Amazon VPC instance is launched, can I change the VPC security groups it belongs to?

- A. Only if the tag "VPC_Change_Group" is true
- B. Yes
- C. You can.
- D. No
- E. You cannot.
- F. Only if the tag "VPC Change Group" is true

Answer: B

NEW QUESTION 298

Select the correct statement:

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

Answer: C

NEW QUESTION 303

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

- A. is not currently
- B. is as of 2013
- C. is planned to be in 2014
- D. will never be

Answer: A

NEW QUESTION 306

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

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

Answer: A

NEW QUESTION 307

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

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

Answer: D

NEW QUESTION 311

When should I choose Provisioned IOPS over Standard RDS storage?

- A. If you use production online transaction processing (OLTP) workloads.
- B. If you have batch-oriented workloads
- C. If you have workloads that are not sensitive to consistent performance

Answer: A

NEW QUESTION 314

In the 'Detailed' monitoring data available for your Amazon EBS volumes, Provisioned IOPS volumes automatically send _ minute metrics to Amazon CloudWatch.

- A. 5
- B. 2
- C. 1
- D. 3

Answer: C

NEW QUESTION 315

Can the string value of 'Key' be prefixed with laws'?

- A. No
- B. Only for EC2 not 53
- C. Yes
- D. Only for 53 not EC

Answer: A

NEW QUESTION 319

What does Amazon RDS stand for?

- A. Regional Data Sewer.
- B. Relational Database Service.
- C. Nothing.
- D. Regional Database Service.

Answer: B

NEW QUESTION 320

What is the maximum response time for a Business level Premium Support case?

- A. 30 minutes
- B. You always get instant responses (within a few seconds).
- C. 10 minutes
- D. 1 hour

Answer: D

NEW QUESTION 321

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

- A. DB Instance
- B. DB Unit
- C. DB Sen/er
- D. DB Volume

Answer: A

NEW QUESTION 324

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

- A. Yes, You can.
- B. N
- C. You cannot

D. Depends on the state of the interface at the time

Answer: B

NEW QUESTION 328

Within the IAM service a GROUP is regarded as a:

- A. A collection of AWS accounts
- B. It's the group of EC2 machines that gain the permissions specified in the GROUP.
- C. There's no GROUP in IAM, but only USERS and RESOURCES.
- D. A collection of user

Answer: D

NEW QUESTION 329

To help you manage your Amazon EC2 instances, images, and other Amazon EC2 resources, you can assign your own metadata to each resource in the form of

--

- A. special filters
- B. functions
- C. tags
- D. wildcards

Answer: C

NEW QUESTION 330

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

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

Answer: A

NEW QUESTION 335

Will I be alerted when automatic fail over occurs?

- A. Only if SNS configured
- B. No
- C. Yes
- D. Only if Cloudwatch configured

Answer: C

NEW QUESTION 338

How can an EBS volume that is currently attached to an EC2 instance be migrated from one Availability Zone to another?

- A. Detach the volume and attach it to another EC2 instance in the other AZ.
- B. Simply create a new volume in the other AZ and specify the original volume as the source.
- C. Create a snapshot of the volume, and create a new volume from the snapshot in the other AZ.
- D. Detach the volume, then use the `ec2-migrate-volume` command to move it to another AZ.

Answer: C

NEW QUESTION 340

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

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

Answer: C

NEW QUESTION 341

Your application provides data transformation services. Files containing data to be transformed are first uploaded to Amazon S3 and then transformed by a fleet of spot EC2 instances. Files submitted by your premium customers must be transformed with the highest priority. How should you implement such a system?

- A. Use a DynamoDB table with an attribute defining the priority level.
- B. Transformation instances will scan the table for tasks, sorting the results by priority level.
- C. Use Route 53 latency based-routing to send high priority tasks to the closest transformation instances.
- D. Use two SQS queues, one for high priority messages, the other for default priority.

- E. Transformation instances first poll the high priority queue; if there is no message, they poll the default priority queue.
- F. Use a single SQS queue
- G. Each message contains the priority level
- H. Transformation instances poll high-priority messages first

Answer: C

NEW QUESTION 346

Which technique can be used to integrate AWS IAM (Identity and Access Management) with an on-premise LDAP (Lightweight Directory Access Protocol) directory service?

- A. Use an IAM policy that references the LDAP account identifiers and the AWS credentials.
- B. Use SAML (Security Assertion Markup Language) to enable single sign-on between AWS and LDAP.
- C. Use AWS Security Token Service from an identity broker to issue short-lived AWS credentials.
- D. Use IAM roles to automatically rotate the IAM credentials when LDAP credentials are updated.
- E. Use the LDAP credentials to restrict a group of users from launching specific EC2 instance type

Answer: B

NEW QUESTION 348

A customer is leveraging Amazon Simple Storage Service in eu-west-1 to store static content for a web-based property. The customer is storing objects using the Standard Storage class. Where are the customer's objects replicated?

- A. A single facility in eu-west-1 and a single facility in eu-central-1
- B. A single facility in eu-west-1 and a single facility in us-east-1
- C. Multiple facilities in eu-west-1
- D. A single facility in eu-west-1

Answer: C

NEW QUESTION 349

In AWS, which security aspects are the customer's responsibility? Choose 4 answers

- A. Security Group and ACL (Access Control List) settings
- B. Decommissioning storage devices
- C. Patch management on the EC2 instance's operating system
- D. Life-cycle management of IAM credentials
- E. Controlling physical access to compute resources
- F. Encryption of EBS (Elastic Block Storage) volumes

Answer: ACDF

Explanation:

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

NEW QUESTION 350

You have a web application running on six Amazon EC2 instances, consuming about 45% of resources on each instance. You are using auto-scaling to make sure that six instances are running at all times.

The number of requests this application processes is consistent and does not experience spikes. The application is critical to your business and you want high availability at all times. You want the load to be distributed evenly between all instances. You also want to use the same Amazon Machine Image (AMI) for all instances. Which of the following architectural choices should you make?

- A. Deploy 6 EC2 instances in one availability zone and use Amazon Elastic Load Balancer.
- B. Deploy 3 EC2 instances in one region and 3 in another region and use Amazon Elastic Load Balancer.
- C. Deploy 3 EC2 instances in one availability zone and 3 in another availability zone and use Amazon Elastic Load Balancer.
- D. Deploy 2 EC2 instances in three regions and use Amazon Elastic Load Balancer

Answer: C

Explanation:

Reference: https://media.amazonwebservices.com/pdf/AWS_Security_Whitepaper.pdf (page 8)

NEW QUESTION 354

Which of the following items are required to allow an application deployed on an EC2 instance to write data to a DynamoDB table? Assume that no security keys are allowed to be stored on the EC2 instance. (Choose 2 answers)

- A. Create an IAM Role that allows write access to the DynamoDB table.
- B. Add an IAM Role to a running EC2 instance.
- C. Create an IAM User that allows write access to the DynamoDB table.
- D. Add an IAM User to a running EC2 instance.
- E. Launch an EC2 Instance with the IAM Role included in the launch configuration

Answer: AE

Explanation:

Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/TicTacToe.Phase3.html>

NEW QUESTION 356

You are working with a customer who has 10 TB of archival data that they want to migrate to Amazon Glacier. The customer has a 1-Mbps connection to the Internet. Which service or feature provides the fastest method of getting the data into Amazon Glacier?

- A. Amazon Glacier multipart upload
- B. AWS Storage Gateway
- C. VM Import/Export
- D. AWS Import/Export

Answer: A

Explanation:

<http://docs.aws.amazon.com/amazonglacier/latest/dev/uploading-archive-mpu.html>

NEW QUESTION 360

How can you secure data at rest on an EBS volume?

- A. Attach the volume to an instance using EC2's SSL interface.
- B. Write the data randomly instead of sequentially.
- C. Encrypt the volume using the 53 server-side encryption service.
- D. Create an IAM policy that restricts read and write access to the volume.
- E. Use an encrypted file system on top of the EBS volum

Answer: E

Explanation:

Reference: <https://aws.amazon.com/blogs/aws/protect-your-data-with-new-ebs-encryption/>

NEW QUESTION 364

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

- A. Enable AWS CloudTrail for the load balancer.
- B. Enable access logs on the load balancer.
- C. Install the Amazon CloudWatch Logs agent on the load balancer.
- D. Enable Amazon CloudWatch metrics on the load balance

Answer: A

NEW QUESTION 366

Which of the following are valid statements about Amazon S3? Choose 2 answers

- A. S3 provides read-after-write consistency for any type of PUT or DELETE
- B. Consistency is not guaranteed for any type of PUT or DELETE
- C. A successful response to a PUT request only occurs when a complete object is saved.
- D. Partially saved objects are immediately readable with a GET after an overwrite PUT.
- E. S3 provides eventual consistency for overwrite PUTS and DELETE

Answer: CE

Explanation:

Reference:

<http://api-portal.anypoint.mulesoft.com/amazon/api/amazon-s3-api/docs/concepts#DataConsistencyNlod el>

NEW QUESTION 370

Which Amazon Elastic Compute Cloud feature can you query from within the instance to access instance properties?

- A. Instance user data
- B. Resource tags
- C. Instance metadata
- D. Amazon Machine Image

Answer: C

NEW QUESTION 371

A customer needs corporate IT governance and cost oversight of all AWS resources consumed by its dMsions. The dMsions want to maintain administrative control of the discrete AWS resources they consume and keep those resources separate from the resources of other dMsions. Which of the following options, when used together will support the autonomy/control of dMsions while enabling corporate IT to maintain governance and cost oversight? Choose 2 answers

- A. Use AWS Consolidated Billing and disable AWS root account access for the child accounts.
- B. Enable IAM cross-account access for all corporate IT administrators in each child account.
- C. Create separate VPCs for each dMsion within the corporate IT AWS account.
- D. Use AWS Consolidated Billing to link the dMsions' accounts to a parent corporate account.
- E. Write all child AWS CloudTrail and Amazon CloudWatch logs to each child account's Amazon S3 'Log'bucket

Answer: DE

NEW QUESTION 373

When an EC2 instance that is backed by an 53-based AMI is terminated, what happens to the data on the root volume?

- A. Data is automatically saved as an EBS snapshot.
- B. Data is automatically saved as an EBS volume.
- C. Data is unavailable until the instance is restarted.
- D. Data is automatically deleted.

Answer: D

NEW QUESTION 377

A company needs to monitor the read and write IOPs metrics for their AWS MySQL RDS instance and send real-time alerts to their operations team. Which AWS services can accomplish this? Choose 2 answers

- A. Amazon Simple Email Service
- B. Amazon CloudWatch
- C. Amazon Simple Queue Service
- D. Amazon Route 53
- E. Amazon Simple Notification Service

Answer: BE

NEW QUESTION 378

.....

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>