

AWS_CLF-C01 Exam

Volume: 325 Questions

Question No:1

Which statement is true about AWS Config and Regions?

- A. AWS Config can only be used in one Region at a time
- B. AWS Config is a Global service, once enabled it will work across all supported regions automatically
- C. AWS Config is not enabled at a Region level
- D. AWS Config is a Region specific service, meaning it has to be configured in every region you wish to use it

Answer: D

Explanation: AWS Config is Region specific, meaning that if you have resources in multiple regions then you will have to configure AWS Config for each Region you want to record resource changes. When doing so, you are able to specify different options for each Region. For example, you could configure Config in one Region to record all supported resources across all services within that Region, and add a predefined AWS Managed Config rule that will check if EBS volumes are encrypted. In another Region you could select to only record a specific type of resource, such as Security Groups with no predefined rules allocated.

Reference: <http://docs.aws.amazon.com/config/latest/developerguide/gs-console.html>

Question No:2

AWS uses the term _____ to describe the ability to scale computing resources up and down easily, with minimal friction.

- A. Scalable
- B. Flexible
- C. Elastic
- D. Large-scale

Answer: C

Explanation: Amazon AWS is based on the concept of elasticity: it means that you can use

AWS_CLF-C01 Exam

instances and other AWS resources without any restrictions on scalability or limits on the amount of available resources.

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

Question No:3

When does the billing process for an Amazon EC2 system begin?

- A. It starts when the Amazon EC2 transitions to the running state.
- B. It starts when the Status column for your distribution changes from Creating to Deployed.
- C. It starts when your instance reaches 720 instance hours.
- D. It starts as soon as you click the create instance option on the main EC2 console.

Answer: A

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

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

Question No:4

Which other AWS service can you use to enable greater security of your CloudTrail log files?

- A. Key Management Service (KMS)
- B. Simple Notification Service (SNS)
- C. CloudWatch
- D. Simple Storage Service (S3)

Answer: A

Explanation: The use of AWS KMS is an optional element of CloudTrail, but it allows additional encryption to be added to your Log files when stored on S3

Reference: <http://docs.aws.amazon.com/kms/latest/developerguide/services-cloudtrail.html>

AWS_CLF-C01 Exam

Question No:5

Your customers are concerned about the security of their sensitive data and their inquiry asks about what happens to old storage devices on AWS. What would be the best answer to this question?

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

Answer: A

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

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

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

Reference: <https://d0.awsstatic.com/whitepapers/aws-security-whitepaper.pdf>

Question No:6

You have been handed a new scope of work from your manager which involves migrating an on-premise application architecture to AWS. During your design process you have to give consideration to current on-premise security and determine which security attributes you are responsible for on AWS and which is the responsibility of AWS. Which of the following is the responsibility of AWS as part of the shared responsibility model? (Choose 3 answers)

- A. Virtualization infrastructure
- B. Edge Locations
- C. Data Center Hardware
- D. Client side Encryption

AWS_CLF-C01 Exam

Answer: A,B,C

Explanation: While AWS manages security of the cloud, security in the cloud is the responsibility of the customer. Customers retain control of what security they choose to implement to protect their own content, platform, applications, systems and networks, no differently than they would for applications in an on-site datacenter.

Reference: <https://aws.amazon.com/compliance/shared-responsibility-model/>

Question No:7

_____ is the process in which a system you have authenticated to establishes what you can access and at what level

- A. Verification
- B. Authorisation
- C. Authentication
- D. Access Control

Answer: B

Explanation: So here we are really looking at your access privileges and permissions. If for example you logged into an AWS account, you would have authenticated yourself by providing the correct identity and password, now AWS security features, and in this case AWS IAM, (Identity & Access Management service), defines the level of authorised access assigned to that identity within the AWS environment. Each identity can have a different level of authorisation properties associated to it. It's these properties that determine what that identity can then access.

Question No:8

In regards to AWS infrastructure, a(n) _____ is deployed in highly populated areas to cache data and reduce latency.

- A. availability zone
- B. edge location
- C. data center
- D. region

[AWS_CLF-C01 Exam](#)

Answer: B

Explanation: Edge Locations are AWS sites deployed in major cities and highly populated areas. Edge Locations are not used to deploy your main infrastructures. Instead they are used by AWS services such as AWS CloudFront to cache data and reduce latency for end user access.

Reference: https://d0.awsstatic.com/whitepapers/AWS_Cloud_Best_Practices.pdf

Question No:9

Which of the following is not a pillar of the Well-Architected Framework developed by AWS?

- A. Security
- B. Performance Efficiency
- C. Profitability
- D. Operational excellence

Answer: C

Explanation: To understand the elements of the well-architected framework you must be aware of the 5 pillars that the framework is based and built upon:

- . Operational excellence
- . Security
- . Reliability
- . Performance efficiency
- . Cost optimization

Reference:

https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf

Question No:10

Which of the following statements best describes Amazon Cognito?

- A. It is an AWS Storage & Content Delivery System
- B. It is a new AWS database
- C. It is a simple user-data synchronization and identity service that helps you securely manage and synchronize app data for your users across their mobile devices.
- D. It is an AWS Deployment & Management System

[AWS_CLF-C01 Exam](#)

Answer: C

Explanation: Amazon Cognito is a simple user-data synchronization and identity service that helps you securely manage and synchronize app data for your users across their mobile devices. You can create unique identifiers for your users through a number of public login providers (Amazon, Facebook, Google or any OpenID Connect compatible provider), or your own user identity system, and also support unauthenticated guests. You can save app data locally on the device allowing your apps to work even when the device is offline and then securely sync and save this data to the AWS cloud, and all of a user's devices.

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

Question No:11

A user has archived data from Amazon S3 to Amazon Glacier. How much data can be restored by the user for free every month?

- A. 5% of archived data
- B. The entire volume of data
- C. 20% of archived data
- D. 15% of archived data

Answer: A

Explanation: When a user has archived an object from S3, the user can restore it only from the S3 console. Glacier charges the user only based on the peak billable rate. The user can restore up to 5% of the archived data free of cost. AWS first calculates the peak billable rate and then based on that calculates the cost of transfer. It never calculates the cost based on the age of the archive.

Reference:

http://aws.amazon.com/s3/faqs/#How_will_I_be_charged_when_restoring_large_amounts_of_data_from_Amazon_Glacier

Question No:12

How is Key Management Service (KMS) priced?

- A. KMS is priced per number of encryption/decryption requests received from all services per month.
- B. KMS is priced per data encryption keys; you are charged for the number of individual data

AWS_CLF-C01 Exam

keys maintained in KMS.

C. KMS is priced per customer master key and the number of requests received per month.

D. KMS is priced per customer master key; you are charged for the number of master keys maintained in KMS.

Answer: C

Explanation: KMS is priced per two factors: the number of Customer Master Keys maintained in KMS and the number of requests received within a month.

Question No:13

What specific type of budgets does AWS budgets not allow you to create?

A. Data transfer

B. Usage

C. Cost

D. Reserved instance utilization

Answer: A

Explanation: Three budget types are available - cost, usage, and reserved instance utilization. Data transfer is not a specific budget type, but the cost related to data transfers would be monitored under the cost budgets.

Reference:

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/budgets-managing-costs.html>

Question No:14

An Amazon Machine Image (AMI) can be best described as a ____.

A. preconfigured template for your instances

B. a virtual machine backup file on a local server hard drive

C. a temporary virtual machine created during horizontal scaling

AWS_CLF-C01 Exam

D. VMware configuration file for any network deployment

Answer: A

Explanation: Amazon EC2 provides Amazon Machine Images (AMIs), which are preconfigured templates for your instances.

Reference: <http://docs.aws.amazon.com/gettingstarted/latest/awsgsg-intro/awsgsg-intro.pdf>

Question No:15

How does AWS define cloud computing?

A. The on-demand delivery of IT resources through a cloud services platform via the Internet with pay-as-you-go pricing.

B. The term used to describe virtualized technology.

C. A physical pool of compute, storage and network resources that can't be accessed over the internet.

D. A pool of servers offering compute resources that are designed to be issued exclusively to individual tenants (users and organizations).

Answer: A

Explanation: Cloud computing provides a simple way to access servers, storage, databases and a broad set of application services over the Internet. A cloud services platform such as Amazon Web Services owns and maintains the network-connected hardware required for these application services, while you provision and use what you need via a web application.

Reference: <https://d0.awsstatic.com/whitepapers/aws-overview.pdf>

Question No:16

What does the phrase 'stop guessing capacity' mean?

A. To set correct data storage lifecycles

B. To use elastic IP addresses to increase high availability

C. Use of auto scaling to prevent the need to predict and guess your capacity and demand requirement

D. To implement self-healing processes

AWS_CLF-C01 Exam

Answer: C

Explanation: One of the best practices of the reliability pillar of the Well-Architected Framework is to 'stop guessing capacity.' This looks at the use of Auto Scaling to prevent the need to predict and guess your capacity and demand requirement which aids in a better end-user experience.

Reference:

https://d1.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf

Question No:17

The function of Key Pairs is to...

- A.encrypt data held on EBS volumes using AES-256 cryptography and then decrypt the data to be read again
- B. safely make programmatic API calls over an encrypted channel
- C.encrypt and decrypt passwords for AWS user accounts
- D. ...encrypt the login information for Linux and Windows EC2 instances, and then decrypt the same information allowing you to authenticate onto the instance.

Answer: D

Explanation: A Key pair as the name implies is made up of 2 components, a Public Key and a Private Key, these keys are 2048-bit SSH-2 RSA keys.

The function of Key Pairs is to encrypt the login information for Linux and Windows EC2 instances, and then decrypt the same information allowing you to authenticate onto the instance.

The Public Key uses Public-Key cryptography to encrypt data such as the username and password. For Window instances the Private Key is used to decrypt this data allowing you to gain access to the login credentials including the password. For Linux instances the Private Key is used to SSH onto the instance.

The Public Key is held and kept by AWS, the Private Key is your responsibility to keep and ensure that it is not lost.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/ec2-key-pairs.html>

Question No:18

Why does AWS recommend using stateless cloud applications, when possible?

- A. They are more always more cost effective than stateful applications.

[AWS_CLF-C01 Exam](#)

- B. Decoupling layers enables more scalability and resilience
- C. They offer a higher degree of encryption for data in transit.
- D. They are more secure than stateful applications.

Answer: B

Explanation: A stateless application needs no knowledge of previous application actions, and stores no session data, and thus can be scaled horizontally with minimal friction.

Reference: https://d0.awsstatic.com/whitepapers/AWS_Cloud_Best_Practices.pdf

Question No:19

What additional method of Access Control can be assigned to an AWS user that utilises a random 6 digit number that is only available for a very short time period before the number changes?

- A. Multi Factor Authentication
- B. Multi Access Authentication
- C. Multi Authorisation Authentication
- D. Varied Factor Authentication

Answer: A

Explanation: IAM allows for Multi-Factor authentication, MFA. This means that any user configured with MFA must use an additional level of authentication as well as a password to be authenticated giving an additional layer of security. This additional authentication utilises a random 6 digit number that is generated by an MFA device that is only available for a very short time period before the number changes again. There is no additional charge for this level of authentication, however you will need your own MFA device, which can be a physical token or a virtual device.

Reference: <https://aws.amazon.com/iam/details/mfa/>

Question No:20

What choice below accurately describes the 'pilot light' disaster recovery method?

- A. Storing critical systems as a template from which resources can be scaled out in the event of a disaster.