TestkingIT

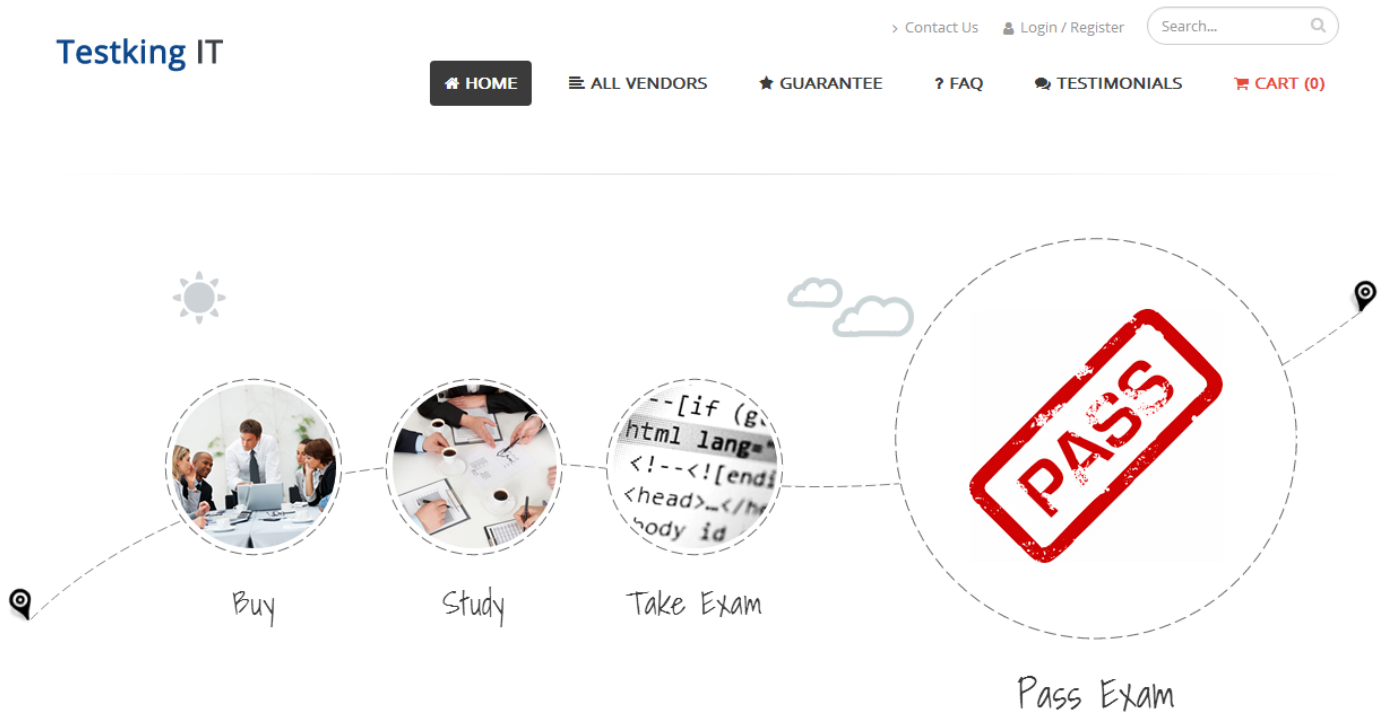

Try **Desktop Test Engine** before you buy

We're not the only ones **happy** about TestKingsIT Practice Material ...

48236+ customers in 100+ countries use TestKingsIT Test Engine. Meet our customers.

iMessenger VOREED GetCustom JET ORANGE iCompany Paradoxx

http://www.testkingit.com/

Latest practice material - Exam Cram - TestKingIT

Exam : **AWS-SysOps**

Title : AWS Certified SysOps Administrator - Associate

Vendor : Amazon

Version : DEMO

NO.1 Can you use the AWS Identity and Access Management (IAM) to assign permissions determining who can manage or modify RDS resources?

A. No, AWS IAM is used only to assign IDs to AWS users.

B. No, AWS IAM is used only to assign activities.

C. No, this permission cannot be assigned by AWS IAM.

D. Yes, you can.

Answer: D

Explanation:

Use AWS Identity and Access Management (IAM) policies to assign permissions that determine who is allowed to manage RDS resources.

For example, you can use IAM to determine who is allowed to create, describe, modify, and delete DB instances, tag resources, or modify DB security groups.

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

NO.2 In regard to AWS CloudFormation, to pass values to your template at runtime you should use _____.

A. resources

B. parameters

C. mapping

D. conditions

Answer: B

Explanation:

Optional parameters are listed in the Parameters section. Parameters enable you to pass values to your template at runtime, and can be dereferenced in the Resources and Outputs sections of the template.

Reference:

http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/concept-template.html

NO.3 A system admin is managing buckets, objects and folders with AWS S3.

Which of the below mentioned statements is true and should be taken in consideration by the sysadmin?

A. The folders support only ACL

B. Both the object and bucket can have an Access Policy but folder cannot have policy

C. Folders can have a policy

D. Both the object and bucket can have ACL but folders cannot have ACL

Answer: D

Explanation:

http://docs.aws.amazon.com/AmazonS3/latest/dev/access-control-overview.html

NO.4 A user has created a VPC with public and private subnets using the VPC wizard.

Which of the below mentioned statements is not true in this scenario?

A. The VPC will create one internet gateway and attach it to VPC

B. The VPC will create two subnets

C. The VPC will launch one NAT instance with an elastic IP

D. The VPC will create a routing instance and attach it with a public subnet

Answer: D

Explanation:

A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create a NAT instance with an elastic IP. Wizard will also create two subnets with route tables. It will also create an internet gateway and attach it to the VPC.

NO.5 True or False: Amazon Route 53 provides highly available and scalable Domain Name System (DNS), domain name registration, and health-checking web services.

A. False, you can only import an existing domain using Amazon Route 53.

B. FALSE

C. TRUE

D. True, however, it only provides .com domains.

Answer: C

Explanation:

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

Reference: http://aws.amazon.com/route53/faqs/

NO.6 A user has setup a billing alarm using CloudWatch for $200.
The usage of AWS exceeded $200 after some days.
The user wants to increase the limit from $200 to $400.
What should the user do?

A. Create a new alarm for the additional $200 amount

B. It is not possible to modify the alarm once it has crossed the usage limit

C. Create a new alarm of $400 and link it with the first alarm

D. Update the alarm to set the limit at $400 instead of $200

Answer: D

Explanation:

AWS CloudWatch supports enabling the billing alarm on the total AWS charges. The estimated charges are calculated and sent several times daily to CloudWatch in the form of metric data.
This data will be stored for 14 days. This data also includes the estimated charges for every service in AWS used by the user, as well as the estimated overall AWS charges. If the user wants to increase the limit, the user can modify the alarm and specify a new threshold.

NO.7 An organization has configured a VPC with an Internet Gateway (IGW). pairs of public and private subnets (each with one subnet per Availability Zone), and an Elastic Load Balancer (ELB) configured to use the public subnets The application s web tier leverages the ELB. Auto Scaling and a mum-AZ RDS database instance The organization would like to eliminate any potential single points ft failure in this design.
What step should you take to achieve this organization's objective?

A. Create and configure a second Elastic Load Balancer to provide a redundant load balancer.

B. Nothing, there are no single points of failure in this architecture.

C. Create and attach a second IGW to provide redundant internet connectivity.
D. Create a second multi-AZ RDS instance in another Availability Zone and configurereplication to provide a redundant database.
Answer: B
Explanation:
is designed to be HA across Availability Zones. You need multiple ELB if you want HA across regions.
"AWS Load Balancer - Cross Network
Many times it happens that after setting up your ELB, you experience significant drops in your performance. The best way to handle this situation is to start with identifying whether your ELB is single AZ or multiple AZ, as single AZ ELB is also considered as one of the Single Points of Failures on AWS Cloud. Once you identify your ELB, it is necessary to make sure ELB loads are kept cross regions."
https://www.botmetric.com/blog/eliminating-single-points-of-failures-on-aws-cloud/

NO.8 A Content Processing team has notified a SysOps Administrator that their content is sometimes taking a long time to process, whereas other times it processes quickly. The Content Processing submits messages to an Amazon Simple Queue Service (Amazon SQS) queue, which details the files that need to be processed. An Amazon EC2 instance polls the queue to determine which file to process next.
How could the Administrator maintain a fast but cost-effective processing time?
A. Attach an Auto Scaling policy to the Amazon SQS queue to increase the number of EC2 instances based on the depth of the SQS queue
B. Create an Auto Scaling policy to increase the number of EC2 instances polling the queue and a CloudWatch alarm to scale based on ApproximateNumberOfMessagesVisible
C. Create an Auto Scaling policy to increase the number of EC2 instances polling the queue and a CloudWatch alarm to scale based on MaxVisibility Timeout
D. Attach an Auto Scaling policy to the SQS queue to scale instances based on the depth of the dead-letter queue
Answer: B
Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-using-sqs-queue.html

NO.9 A user has setup an Auto Scaling group.
The group has failed to launch a single instance for more than 24 hours.
What will happen to Auto Scaling in this condition?
A. Auto Scaling will start an instance in a separate region
B. Auto Scaling will suspend the scaling process
C. Auto Scaling will keep trying to launch the instance for 72 hours
D. The Auto Scaling group will be terminated automatically
Answer: B
Explanation:
If Auto Scaling is trying to launch an instance and if the launching of the instance fails continuously, it will suspend the processes for the Auto Scaling groups since it repeatedly failed to launch an instance. This is known as an administrative suspension. It commonly applies to the Auto Scaling group that has no running instances which is trying to launch instances for more than 24 hours, and

has not succeeded in that to do so.

NO.10 A SysOps Administrator must ensure that AWS CloudFormation deployment changes are properly tracked for governance.
Which AWS service should be used to accomplish this?
A. Amazon Inspector
B. AWS Trusted Advisor
C. AWS Artifact
D. AWS Config
Answer: D

NO.11 A user has created a queue named "myqueue" with SQS.
There are four messages published to queue which are not received by the consumer yet.
If the user tries to delete the queue, what will happen?
A. It will delete the queue
B. It will initiate the delete but wait for four days before deleting until all messages are deleted automatically.
C. A user can never delete a queue manually.
AWS deletes it after 30 days of inactivity on queue
D. I t will ask user to delete the messages first
Answer: A
Explanation:
SQS allows the user to move data between distributed components of applications so they can perform different tasks without losing messages or requiring each component to be always available. The user can delete a queue at any time, whether it is empty or not. It is important to note that queues retain messages for a set period of time. By default, a queue retains messages for four days.

NO.12 A photo-sharing service stores pictures in Amazon Simple Storage Service (S3) and allows application sign-in using an OpenID Connect-compatible identity provider. Which AWS Security Token Service approach to temporary access should you use for the Amazon S3 operations?
A. Cross-Account Access
B. Web Identity Federation
C. AWS Identity and Access Management roles
D. SAML-based Identity Federation
Answer: B

NO.13 An Administrator has an Amazon EC2 instance with an IPv6 address. The Administrator needs to prevent direct access to this instance from the Internet.
The Administrator should place the EC2 instance in a:
A. Private Subnet with an egress-only Internet Gateway attached to the subnet and placed in the subnet Route Table.
B. Public subnet and a security group that blocks inbound IPv6 traffic attached to the interface.
C. Public subnet with an egress-only Internet Gateway attached to the VPC and placed in the VPC Route Table.

D. Private subnet with an egress-only Internet Gateway attached to the VPC and placed in the subnet Route Table.

Answer: B

Explanation:

Any IPv6 IPs are public. So does not matter where do you put the instance, the only way to block inbound traffic is to use Security Groups.

NO.14 The Accounting department would like to receive billing updates more than once a month. They would like the updates to be in a format that can easily be viewed with a spreadsheet application.

How can this request be fulfilled?

A. Use AWS Lambda, triggered by CloudWatch, to query billing data and push to Amazon RDS.

B. Use Amazon CloudWatch Events to schedule a billing inquiry on a bi-weekly basis. Use AWS Glue to convert the output to CSV.

C. Use the AWS CLI to output billing data as JSON. Use Amazon SES to email bills on a daily basis.

D. Set AWS Cost and Usage Reports to publish bills daily to an Amazon S3 bucket in CSV format.

Answer: D

NO.15 An organization is running multiple applications for their customers. Each application is deployed by running a base AWS CloudFormation template that configures a new VPC. All applications are run in the same AWS account and AWS Region. A SysOps Administrator has noticed that when trying to deploy the same AWS CloudFormation stack, it fails to deploy.

What is likely to be the problem?

A. The account has reached the default limit for VPCs allowed.

B. The VPC configuration parameters have changed and must be updated in the template.

C. The AWS CloudFormation template needs to be updated to the latest version.

D. The Amazon Machine image used is not available in that region.

Answer: A

Explanation:

The default VPC Limitation per region is 5.

NO.16 You are setting up a VPC and you need to set up a public subnet within that VPC. Which following requirement must be met for this subnet to be considered a public subnet?

A. Subnet's traffic is not routed to an internet gateway but has its traffic routed to a virtual private gateway.

B. Subnet's traffic is routed to an internet gateway.

C. Subnet's traffic is not routed to an internet gateway.

D. None of these answers can be considered a public subnet.

Answer: B

Explanation:

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

A subnet is a range of IP addresses in your VPC. You can launch AWS resources into a subnet that you select. Use a public subnet for resources that must be connected to the internet, and a private subnet for resources that won't be connected to the Internet. If a subnet's traffic is routed to an internet gateway, the subnet is known as a public subnet. If a subnet doesn't have a route to the internet gateway, the subnet is known as a private subnet. If a subnet doesn't have a route to the internet gateway, but has its traffic routed to a virtual private gateway, the subnet is known as a VPN-only subnet.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

NO.17 A company backs up data from its data center using a tape gateway on AWS Storage Gateway.
The SysOps Administrator needs to reboot the virtual machine running Storage Gateway.
What process will protect data integrity?
A. Stop Storage Gateway and reboot the virtual machine, then restart Storage Gateway.
B. Shut down the virtual machine and stop Storage Gateway, then turn on the virtual machine.
C. Reboot the virtual machine, then restart Storage Gateway.
D. Reboot the virtual machine.
Answer: A

NO.18 Amazon S3 provides a number of security features for protection of data at rest, which you can use or not, depending on your threat profile. What feature of S3 allows you to create and manage your own encryption keys for sending data?
A. Client-side Encryption
B. Data integrity compromise
C. Network traffic protection
D. Server-side Encryption
Answer: A
Explanation:
With client-side encryption you create and manage your own encryption keys. Keys you create are not exported to AWS in clear text. Your applications encrypt data before submitting it to Amazon S3, and decrypt data after receiving it from Amazon S3. Data is stored in an encrypted form, with keys and algorithms only known to you. While you can use any encryption algorithm, and either symmetric or asymmetric keys to encrypt the data, the AWS-provided Java SDK offers Amazon S3 client-side encryption features.
Reference: https://d0.awsstatic.com/whitepapers/aws-security-best-practices.pdf

NO.19 Which of the following statements is true of IAM?
A. None of these are correct.
B. If you are configuring MFA for a user who will use a smartphone to generate an OTP, you must have the smartphone available in order to finish the wizard.
C. If you are configuring MFA for a user who will use a smartphone to generate an OTP, you can finish the wizard on any device and later use the smartphone for authentication.
D. If you are configuring MFA for a user who will use a smartphone to generate an OTP, the smartphone is not required in order to finish the wizard.

Answer: B
Explanation:
MFA can be used either with a specific MFA-enabled device or by installing an application on a smartphone. If a user chooses to use her smartphone, physical access to the device is required in order to complete the configuration wizard.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/GenerateMFAConfig.html

NO.20 In AWS, which security aspects are the customer's responsibility? (Choose four.)
A. Life-cycle management of IAM credentials
B. Patch management on the EC2 instance s operating system
C. Security Group and ACL (Access Control List) settings
D. Decommissioning storage devices
E. Controlling physical access to compute resources
F. Encryption of EBS (Elastic Block Storage) volumes
Answer: A,B,C,F

NO.21 George has shared an EC2 AMI created in the US East region from his AWS account with Stefano.
George copies the same AMI to the US West region.
Can Stefano access the copied AMI of George's account from the US West region?
A. It is not possible to share the AMI with a specific account
B. Yes, since copy AMI copies all private account sharing permissions
C. No, copy AMI does not copy the permission
D. Yes, since copy AMI copies all the permissions attached with the AMI
Answer: C
Explanation:
Within EC2, when the user copies an AMI, the new AMI is fully independent of the source AMI; there is no link to the original (source. AMI. AWS does not copy launch the permissions, user- defined tags or the Amazon S3 bucket permissions from the source AMI to the new AMI.
Thus, in this case by default Stefano will not have access to the AMI in the US West region.

NO.22 A company application stores document within an Amazon S3 bucket. The application is running on Amazon EC3 in a VPC. A recent change in security requirement states traffic between the company's application and the S3 bucket must leave the Amazon network.
What AWS feature can provide this functionality?
A. Security groups
B. NAT gateways
C. Gateway VPC endpoint
D. Virtual private gateway
Answer: C
Explanation:
A VPC endpoint enables you to create a private connection between your VPC and another AWS service without requiring access over the Internet, through a NAT device, a VPN connection, or AWS Direct Connect. Endpoints are virtual devices.

NO.23 Is it possible to access S3 objects from the Internet?
A. Yes, but it has to pass through EC2.
B. Yes, it is possible if proper public readable accesses and ACLs are set.
C. No, only a general overview of S3 objects can be read from the Internet.
D. No, there is no way to access any S3 objects from the Internet.
Answer: B
Explanation:
You must grant read permission on the specific objects to make them publicly accessible so that your users can view them on your website. You make objects publicly readable by using either the object ACL or by writing a bucket policy.
Reference: https://aws.amazon.com/articles/5050

NO.24 A placement group in Amazon EC2 can
A. isolate any instance-type physically so that groups access local resources.
B. logically name and tag different tiers of the system (DB, application, business logic etc).
C. reduce network latency and increase network throughput
D. place high memory instances in one logical group.
Answer: C
Explanation:
A placement group is a logical grouping of instances within a single Availability Zone. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. To provide the lowest latency, and the highest packet-per-second network performance for your placement group, choose an instance type that supports enhanced networking.
Reference: https://aws.amazon.com/ec2/faqs/

NO.25 A company uses AWS CloudFormation to deploy its application infrastructure. Recently, a user accidentally changed a property of a database in a CloudFormation template and performed a stack update that caused an interruption to the application. A SysOps Administrator must determine how to modify the deployment process to allow the DevOps team to continue to deploy the infrastructure, but prevent against accidental modifications to specific resources.
Which solution will meet these requirements?
A. Launch the CloudFormation templates using a stack policy with an explicit allow for all resources and an explicit deny of the protected resources with an action of Update:*
B. Attach an IAM policy to the DevOps team role that prevents a CloudFormation stack from updating, with a condition based on the specific Amazon Resource names (ARNs) of the protected resources.
C. Set up an Amazon CloudWatch Events event with a rule to trigger based on any CloudFormation API call. An AWS Lambda function can then describe the stack to determine if any protected resources were modified and cancel the operation.
D. Set up an AWS Config rule to alert based on changes to any Cloud Formation stack. An AWS Lambda function can then describe the stack to determine if any protected resources were modified and cancel the operation.
Answer: A

NO.26 A system admin is planning to encrypt all objects being uploaded to S3 from an application. The system admin does not want to implement his own encryption algorithm; instead he is planning to use server side encryption by supplying his own key (SSE-C..
Which parameter is not required while making a call for SSE-C?

A. x-amz-server-side-encryption-customer-key-MD5

B. x-amz-server-side-encryption-customer-algorithm

C. x-amz-server-side-encryption-customer-key-AES-256

D. x-amz-server-side-encryption-customer-key

Answer: C

Explanation:
AWS S3 supports client side or server side encryption to encrypt all data at rest. The server side encryption can either have the S3 supplied AES-256 encryption key or the user can send the key along with each API call to supply his own encryption key (SSE-C.. When the user is supplying his own encryption key, the user has to send the below mentioned parameters as a part of the API calls:
x-amz-server-side-encryption-customer-algorithm:
Specifies the encryption algorithm x-amz-server-side-encryption-customer-key:
To provide the base64-encoded encryption key
x-amz-server-side-encryption-customer-key-MD5:
To provide the base64-encoded 128-bit MD5 digest of the encryption key

NO.27 A web application runs on Amazon EC2 instances with public IPs assigned behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones.
The application stores data in an Amazon RDS Multi-AZ DB instance. The Application Load Balancer, EC2 instances, and RDS DB instance all run in separate sets of subnets. The EC2 instances can communicate with the DB instance, but cannot connect with external services.
What is the MOST likely solution?

A. Create and attach a virtual private gateway to the VP Create a route table for the EC2 instances' subnets that sends Internet traffic to the gateway.

B. Create a VPC peering connection to a VPC that has an Internet gateway attached. Create a route table for the EC2 instances' subnets that sends Internet traffic to the peered VPC.

C. Create and attach an Internet gateway to the VPC. Create a route table for the EC2 instance's subnets that sends Internet traffic to the gateway.

D. Assign a public IP address to the database server and restart the database engine.

Answer: C

Explanation:
For internet, we required Internet gateway to attached with vpc and configure route.

NO.28 A company is using an AWS KMS customer master key (CMK) with imported key material. The company references the CMK by its alias in the Java application to encrypt data. The CMK must be rotated every 6 months.
What is the process to rotate the key?

A. Import a copy of the existing key material into a new CMK as a backup, and set the rotation schedule for 6 months.

B. Delete the current key material, and import new material into the existing CMK.

C. Enable automatic key rotation for the CMK, and specify a period of 6 months.
D. Create a new CMK with new imported material, and update the key alias to point to the new CMK.
Answer: D
Explanation:
When you import key material into a CMK, the CMK is permanently associated with that key material. You can reimport the same key material, but you cannot import different key material into that CMK. Also, you cannot enable automatic key rotation for a CMK with imported key material. However, you can manually rotate a CMK with imported key material.
https://docs.aws.amazon.com/kms/latest/developerguide/importing-keys.html

NO.29 An organization has configured Auto Scaling with ELB.
There is a memory issue in the application which is causing CPU utilization to go above 90%.
The higher CPU usage triggers an event for Auto Scaling as per the scaling policy.
If the user wants to find the root cause inside the application without triggering a scaling activity, how can he achieve this?
A. Delete Auto Scaling until research is completed
B. It is not possible to find the root cause from that instance without triggering scaling
C. Suspend the scaling process until research is completed
D. Stop the scaling process until research is completed
Answer: C
Explanation:
Auto Scaling allows the user to suspend and then resume one or more of the Auto Scaling processes in the Auto Scaling group. This is very useful when the user wants to investigate a configuration problem or some other issue, such as a memory leak with the web application and then make changes to the application, without triggering the Auto Scaling process.

NO.30 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 Notification Service
B. Amazon Route 53
C. Amazon CloudWatch
D. Amazon Simple Queue Service
E. Amazon Simple Email Service
Answer: A,C

NO.31 How many metrics are supported by CloudWatch for Auto Scaling ?
A. 1 metric and 5 dimensions
B. 5 metrics and 1 dimension
C. 8 metrics and 1 dimension
D. 7 metrics and 5 dimension
Answer: C
Explanation:
AWS Auto Scaling supports both detailed as well as basic monitoring of the CloudWatch metrics.

Basic monitoring happens every 5 minutes, while detailed monitoring happens every minute. It supports 8 metrics and 1 dimension.
The metrics are:
GroupMinSize
GroupMaxSize
GroupDesiredCapacity
GroupInServiceInstances
GroupPendingInstances
GroupStandbyInstances
GroupTerminatingInstances
GroupTotalInstances
The dimension is AutoScalingGroupName
Reference:
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/supported_services.htm l

NO.32 A company has received a notification in its AWS Personal Health Dashboard that one of its Amazon EBS-backed Amazon EC2 instances is on hardware that is scheduled for maintenance.
The instance runs a critical production workload that must be available during normal business hours.
Which steps will ensure that the instance maintenance does not produce an outage?
A. Create an Amazon Machine Image (AMI) of the instance and use the AMI to launch a new instance once the existing instance is retired.
B. Enable termination protection on the EC2 instance.
C. Stop and start the EC2 instance during a maintenance window outside of normal business hours.
D. Configure an Amazon Lambda function to automatically start the instance if it is stopped.
Answer: D

NO.33 The Statement element, of an AWS IAM policy, contains an array of individual statements. Each individual statement is a(n) _____ block enclosed in braces { }.
A. AJAX
B. JSON
C. jQuery
D. JavaScript
Answer: B
Explanation:
The Statement element, of an IAM policy, contains an array of individual statements. Each individual statement is a JSON block enclosed in braces { }.
Reference:
http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html

NO.34 An organization recently faced a network outage while uploading data into one of their S3 buckets. This outage generated many incomplete multipart uploads in that S3 bucket. A sysops administrator wants to delete the incomplete multipart uploads and ensure that the incomplete multipart uploads are deleted automatically the next time such an event occurs.
How should this be done?
A. Create an Amazon S3 lifecycle rule to abort incomplete multipart uploads so that they are deleted

this time and in the future.

B. Create an Amazon S3 Event Notification to trigger an AWS Lambda function that deletes incomplete multipart uploads.

C. Use the AWS CLI to list all the multipart uploads, and abort all the incomplete uploads from the day of the event so that they are deleted.

D. Use the AWS Management Console to abort all the incomplete uploads from the day of the event so that they are deleted.

Answer: A

Explanation:https://aws.amazon.com/blogs/aws/s3-lifecycle-management-update-support-for-multipart-uploads-and-delete-markers/

NO.35 A user has configured ELB with three instances.
The user wants to achieve High Availability as well as redundancy with ELB.
Which of the below mentioned AWS services helps the user achieve this for ELB?

A. AWS EMR

B. Route 53

C. Auto Scaling

D. AWS Mechanical Turk

Answer: B

Explanation:
The user can provide high availability and redundancy for applications running behind Elastic Load Balancer by enabling the Amazon Route 53 Domain Name System (DNS. failover for the load balancers. Amazon Route 53 is a DNS service that provides reliable routing to the user's infrastructure.

NO.36 A user has created a VPC with the public and private subnets using the VPC wizard.
The VPC has CIDR 20.0.0.0/16.
The public subnet uses CIDR 20.0.1.0/24.
The user is planning to host a web server in the public subnet (port 80. and a DB server in the private subnet (port 3306..
The user is configuring a security group for the public subnet (WebSecGrp. and the private subnet (DBSecGrp..
Which of the below mentioned entries is required in the private subnet database security group (DBSecGrp.?

A. Allow Outbound on port 80 for Destination NAT Instance IP

B. Allow Outbound on port 3306 for Destination Web Server Security Group (WebSecGrp.

C. Allow Inbound on port 3306 from source 20.0.0.0/16

D. Allow Inbound on port 3306 for Source Web Server Security Group (WebSecGrp.

Answer: D

Explanation:
A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet to host the web server and DB server respectively, the user should configure that the instances in the private subnet can receive inbound traffic from the public subnet on the DB port. Thus, configure port 3306 in Inbound with the source as the Web Server Security Group (WebSecGrp.. The user should configure ports 80 and 443 for Destination

0.0.0.0/0 as the route table directs traffic to the NAT instance from the private subnet.

NO.37 Malicious traffic is reaching company web servers from a single IP address located in another country. The SysOps Administrator is tasked with blocking this IP address.
How should the Administrator implement the restriction?
A. Edit the network access control list for the web server subnet and add a deny entry for the IP address
B. Use Amazon CloudFront's geo restriction feature to block traffic from the IP address
C. Edit the security group for the web servers and add a deny entry for the IP address
D. Edit the VPC route table to route the malicious IP address to a black hole
Answer: A
Explanation:
We need to restrict one ip so Geo restriction is false.
A false because you cant deny trafic for one ip in Security Group

NO.38 A company has centralized all its logs into one Amazon CloudWatch Logs log group. The SysOps Administrator is to alert different teams of any issues relevant to them.
What is the MOST efficient approach to accomplish this?
A. Redesign the aggregation of logs so that each team's relevant parts are sent to a separate log group, then subscribe each team to its respective log group.
B. Set up different metric filters for each team based on patterns and alerts. Each alarm will notify the appropriate notification list.
C. Write a AWS lambda function that will query the logs every minute and contain the logic of which team to notify on which patterns and issues.
D. Create an AWS Auto Scaling group of Amazon EC2 instances that will scale based on the amount of ingested log entries. This group will pull streams, look for patterns, and send notifications to relevant teams.
Answer: C
Explanation:
https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html

NO.39 A user is planning to setup notifications on the RDS DB for a snapshot.
Which of the below mentioned event categories is not supported by RDS for this snapshot source type?
A. Backup
B. Deletion
C. Creation
D. Restoration
Answer: A
Explanation:
Amazon RDS uses the Amazon Simple Notification Service to provide a notification when an Amazon RDS event occurs. Event categories for a snapshot source type include: Creation, Deletion, and Restoration. The Backup is a part of DB instance source type.

NO.40 A company is running an Oracle database engine that handles heavy online transaction

processing (OLTP) structured data traffic.
How can a SysOps administrator ensure that the database has high availability?

A. Use Amazon DynamoOB to store the data

B. Use an Amazon Redshift cluster to store the data

C. Use Amazon RDS Multi -AZ deployment to store the data

D. Use Amazon RDS read replicas in a different region to store the data

Answer: C