

## [Full Version 2017 Amazon AWS-SysOps Dumps Free Download 100% Pass Promised By Lead2pass (201-220)]

2017 February Amazon Official New Released AWS-SysOps Dumps in Lead2pass.com! 100% Free Download! 100% Pass Guaranteed! 2017 get prepared with fully updated Amazon AWS-SysOps real exam questions and accurate answers for AWS-SysOps exam. Lead2pass IT experts review the AWS-SysOps newly added questions and offer correct Amazon AWS-SysOps exam questions answers. 100% pass easily! Following questions and answers are all new published by Amazon Official Exam Center: <http://www.lead2pass.com/aws-sysops.html>

**QUESTION 201** An organization is planning to create a user with IAM. They are trying to understand the limitations of IAM so that they can plan accordingly. Which of the below mentioned statements is not true with respect to the limitations of IAM? A. One IAM user can be a part of a maximum of 5 groups B. The organization can create 100 groups per AWS account C. One AWS account can have a maximum of 5000 IAM users D. One AWS account can have 250 roles  
Answer: A Explanation: AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. The default maximums for each of the IAM entities is given below: Groups per AWS account: 100 Users per AWS account: 5000 Roles per AWS account: 250 Number of groups per user: 10 (that is, one user can be part of these many groups.

**QUESTION 202** A user has created a VPC with public and private subnets using the VPC wizard. The user has not launched any instance manually and is trying to delete the VPC. What will happen in this scenario? A. It will not allow to delete the VPC as it has subnets with route tables B. It will not allow to delete the VPC since it has a running route instance C. It will terminate the VPC along with all the instances launched by the wizard D. It will not allow to delete the VPC since it has a running NAT instance  
Answer: D Explanation: A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. 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. If the user is trying to delete the VPC it will not allow as the NAT instance is still running.

**QUESTION 203** An AWS account wants to be part of the consolidated billing of his organization's payee account. How can the owner of that account achieve this? A. The payee account has to request AWS support to link the other accounts with his account B. The owner of the linked account should add the payee account to his master account list from the billing console C. The payee account will send a request to the linked account to be a part of consolidated billing D. The owner of the linked account requests the payee account to add his account to consolidated billing  
Answer: C Explanation: AWS consolidated billing enables the organization to consolidate payments for multiple Amazon Web Services (AWS) accounts within a single organization by making a single paying account. To add a particular account (linked to the master (payee) account, the payee account has to request the linked account to join consolidated billing. Once the linked account accepts the request henceforth all charges incurred by the linked account will be paid by the payee account.

**QUESTION 204** A user is measuring the CPU utilization of a private data centre machine every minute. The machine provides the aggregate of data every hour, such as Sum of data", "Min value", "Max value, and "Number of Data points". The user wants to send these values to CloudWatch. How can the user achieve this? A. Send the data using the put-metric-data command with the aggregate-values parameter B. Send the data using the put-metric-data command with the average-values parameter C. Send the data using the put-metric-data command with the statistic-values parameter D. Send the data using the put-metric-data command with the aggregateata parameter  
Answer: C Explanation: AWS CloudWatch supports the custom metrics. The user can always capture the custom data and upload the data to CloudWatch using CLI or APIs. The user can publish the data to CloudWatch as single data points or as an aggregated set of data points called a statistic set using the command put-metric-data. When sending the aggregate data, the user needs to send it with the parameter statistic-values: awscloudwatch put-metric-data--metric-name <Name>--namespace <Custom namespace>--timestamp <UTC Format>--statistic-values Sum=XX,Minimum=YY,Maximum=AA,SampleCount=BB--unit Milliseconds

**QUESTION 205** A user is planning to set up the Multi AZ feature of RDS. Which of the below mentioned conditions won't take advantage of the Multi AZ feature? A. Availability zone outage B. A manual failover of the DB instance using Reboot with failover option C. Region outage D. When the user changes the DB instance's server type  
Answer: C Explanation: Amazon RDS when enabled with Multi AZ will handle failovers automatically. Thus, the user can resume database operations as quickly as possible without administrative intervention. The primary DB instance switches over automatically to the standby replica if any of the following conditions occur: - An Availability Zone outage - The primary DB instance fails - The DB instance's server type is changed - The DB instance is undergoing software patching - A manual failover of the DB instance was initiated using Reboot with failover

**QUESTION 206** An organization has launched 5 instances: 2 for production and 3 for testing. The organization wants that one particular group of IAM users should only

access the test instances and not the production ones. How can the organization set that as a part of the policy? A. Launch the test and production instances in separate regions and allow region wise access to the group B. Define the IAM policy which allows access based on the instance ID C. Create an IAM policy with a condition which allows access to only small instances D. Define the tags on the test and production servers and add a condition to the IAM policy which allows access to specific tags Answer: D Explanation: AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. The user can add conditions as a part of the IAM policies. The condition can be set on AWS Tags, Time, and Client IP as well as on various parameters. If the organization wants the user to access only specific instances he should define proper tags and add to the IAM policy condition. The sample policy is shown below. "Statement": [ { "Action": "ec2:\*", "Effect": "Allow", "Resource": "\*", "Condition": { "StringEquals": { "ec2:ResourceTag/InstanceType": "Production" } } } ]

QUESTION 207 A user is having data generated randomly based on a certain event. The user wants to upload that data to CloudWatch. It may happen that event may not have data generated for some period due to andomness. Which of the below mentioned options is a recommended option for this case? A. For the period when there is no data, the user should not send the data at all B. For the period when there is no data the user should send a blank value C. For the period when there is no data the user should send the value as 0 D. The user must upload the data to CloudWatch as having no data for some period will cause an error at CloudWatch monitoring Answer: C Explanation: AWS CloudWatch supports the custom metrics. The user can always capture the custom data and upload the data to CloudWatch using CLI or APIs. When the user data is more random and not generated at regular intervals, there can be a period which has no associated data. The user can either publish the zero (0. Value for that period or not publish the data at all. It is recommended that the user should publish zero instead of no value to monitor the health of the application. This is helpful in an alarm as well as in the generation of the sample data count.

QUESTION 208 A user has created a VPC with public and private subnets using the VPC wizard. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.0.0/24 . The NAT instance ID is i-a12345. Which of the below mentioned entries are required in the main route table attached with the private subnet to allow instances to connect with the internet? A. Destination: 0.0.0.0/0 and Target: i-a12345 B. Destination: 20.0.0.0/0 and Target: 80 C. Destination: 20.0.0.0/0 and Target: i-a12345 D. Destination: 20.0.0.0/24 and Target: i-a12345 Answer: A 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 two route tables and attach to the subnets. The main route table will have the entry "Destination: 0.0.0.0/0 and Target: ia12345", which allows all the instances in the private subnet to connect to the internet using NAT.

QUESTION 209 A user has configured ELB with a TCP listener at ELB as well as on the back-end instances. The user wants to enable a proxy protocol to capture the source and destination IP information in the header. Which of the below mentioned statements helps the user understand a proxy protocol with TCP configuration? A. If the end user is requesting behind a proxy server then the user should not enable a proxy protocol on ELB B. ELB does not support a proxy protocol when it is listening on both the load balancer and the back-end instances C. Whether the end user is requesting from a proxy server or directly, it does not make a difference for the proxy protocol D. If the end user is requesting behind the proxy then the user should add the "isproxy" flag to the ELB Configuration Answer: A Explanation: When the user has configured Transmission Control Protocol (TCP. or Secure Sockets Layer (SSL. for both front-end and back-end connections of the Elastic Load Balancer, the load balancer forwards the request to the back-end instances without modifying the request headers unless the proxy header is enabled. If the end user is requesting from a Proxy Protocol enabled proxy server, then the ELB admin should not enable the Proxy Protocol on the load balancer. If the Proxy Protocol is enabled on both the proxy server and the load balancer, the load balancer will add another header to the request which already has a header from the proxy server. This duplication may result in errors.

QUESTION 210 A user has created a queue named "awsmodule" with SQS. One of the consumers of queue is down for 3 days and then becomes available. Will that component receive message from queue? A. Yes, since SQS by default stores message for 4 days B. No, since SQS by default stores message for 1 day only C. No, since SQS sends message to consumers who are available that time D. Yes, since SQS will not delete message until it is delivered to all consumers 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. Queues retain messages for a set period of time. By default, a queue retains messages for four days. However, the user can configure a queue to retain messages for up to 14 days after the message has been sent.

QUESTION 211 A user has configured an ELB to distribute the traffic among multiple instances. The user instances are facing some issues due to the back-end servers. Which of the below mentioned CloudWatch metrics helps the user understand the issue with the instances? A. HTTPCode\_Backend\_3XX B. HTTPCode\_Backend\_4XX C. HTTPCode\_Backend\_2XX D. HTTPCode\_Backend\_5XX Answer: D Explanation: CloudWatch is used to monitor AWS as well as the custom services. For ELB,

CloudWatch provides various metrics including error code by ELB as well as by back-end servers (instances.. It gives data for the count of the number of HTTP response codes generated by the back-end instances. This metric does not include any response codes generated by the load balancer. These metrics are: The 2XX class status codes represents successful actions The 3XX class status code indicates that the user agent requires action The 4XX class status code represents client errors The 5XX class status code represents back-end server errors

QUESTION 212 A user is displaying the CPU utilization, and Network in and Network out CloudWatch metrics data of a single instance on the same graph. The graph uses one Y-axis for CPU utilization and Network in and another Y-axis for Network out. Since Network in is too high, the CPU utilization data is not visible clearly on graph to the user. How can the data be viewed better on the same graph? A. It is not possible to show multiple metrics with the different units on the same graph B. Add a third Y-axis with the console to show all the data in proportion C. Change the axis of Network by using the Switch command from the graph D. Change the units of CPU utilization so it can be shown in proportion with Network

Answer: C

Explanation: Amazon CloudWatch provides the functionality to graph the metric data generated either by the AWS services or the custom metric to make it easier for the user to analyse. It is possible to show the multiple metrics with different units on the same graph. If the graph is not plotted properly due to a difference in the unit data over two metrics, the user can change the Y-axis of one of the graph by selecting that graph and clicking on the Switch option.

QUESTION 213 A user has created an Auto Scaling group with default configurations from CLI. The user wants to setup the CloudWatch alarm on the EC2 instances, which are launched by the Auto Scaling group. The user has setup an alarm to monitor the CPU utilization every minute. Which of the below mentioned statements is true? A. It will fetch the data at every minute but the four data points [corresponding to 4 minutes] will not have value since the EC2 basic monitoring metrics are collected every five minutes B. It will fetch the data at every minute as detailed monitoring on EC2 will be enabled by the default launch configuration of Auto Scaling C. The alarm creation will fail since the user has not enabled detailed monitoring on the EC2 instances D. The user has to first enable detailed monitoring on the EC2 instances to support alarm monitoring at every minute

Answer: B

Explanation: CloudWatch is used to monitor AWS as well as the custom services. To enable detailed instance monitoring for a new Auto Scaling group, the user does not need to take any extra steps. When the user creates an Auto Scaling launch config using CLI, each launch configuration contains a flag named InstanceMonitoring.Enabled. The default value of this flag is true. Thus, by default detailed monitoring will be enabled for Auto Scaling as well as for all the instances launched by that Auto Scaling group.

QUESTION 214 A root account owner has given full access of his S3 bucket to one of the IAM users using the bucket ACL. When the IAM user logs in to the S3 console, which actions can he perform? A. He can just view the content of the bucket B. He can do all the operations on the bucket C. It is not possible to give access to an IAM user using ACL D. The IAM user can perform all operations on the bucket using only API/SDK

Answer: C

Explanation: Each AWS S3 bucket and object has an ACL (Access Control List. associated with it. An ACL is a list of grants identifying the grantee and the permission granted. The user can use ACLs to grant basic read/write permissions to other AWS accounts. ACLs use an Amazon S3?pecific XML schema. The user cannot grant permissions to other users (IAM users. in his account.

QUESTION 215 A user is planning to scale up an application by 8 AM and scale down by 7 PM daily using Auto Scaling. What should the user do in this case? A. Setup the scaling policy to scale up and down based on the CloudWatch alarms B. The user should increase the desired capacity at 8 AM and decrease it by 7 PM manually C. The user should setup a batch process which launches the EC2 instance at a specific time D. Setup scheduled actions to scale up or down at a specific time

Answer: D

Explanation: Scale based on a schedule Sometimes you know exactly when you will need to increase or decrease the number of instances in your group, simply because that need arises on a predictable schedule. Scaling by schedule means that scaling actions are performed automatically as a function of time and date. For more information, see Scheduled Scaling.

[http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/schedule\\_time.html](http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/schedule_time.html)

QUESTION 216 A user has configured Auto Scaling with the minimum capacity as 2 and the desired capacity as 2. The user is trying to terminate one of the existing instance with the command: `as-terminate-instance-in-auto-scaling-group<Instance ID>--decrement-desired-capacity` What will Auto Scaling do in this scenario? A. Terminates the instance and does not launch a new instance B. Terminates the instance and updates the desired capacity to 1 C. Terminates the instance and updates the desired capacity and minimum size to 1 D. Throws an error

Answer: D

Explanation: The Auto Scaling command `as-terminate-instance-in-auto-scaling-group <Instance ID>` will terminate the specific instance ID. The user is required to specify the parameter `as--decrement-desired-capacity`. Then Auto Scaling will terminate the instance and decrease the desired capacity by 1. In this case since the minimum size is 2, Auto Scaling will not allow the desired capacity to go below 2. Thus, it will throw an error.

QUESTION 217 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. No, copy AMI does not copy the permission B. It is not possible to share the AMI with a specific account C. Yes, since copy AMI copies all private account sharing permissions D. Yes, since copy AMI

copies all the permissions attached with the AMI Answer: A 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. QUESTION 218 A user has moved an object to Glacier using the life cycle rules. The user requests to restore the archive after 6 months. When the restore request is completed the user accesses that archive. Which of the below mentioned statements is not true in this condition? A. The archive will be available as an object for the duration specified by the user during the restoration request B. The restored object's storage class will be RRS C. The user can modify the restoration period only by issuing a new restore request with the updated period D. The user needs to pay storage for both RRS (restored. and Glacier (Archive. Rates Answer: B Explanation: AWS Glacier is an archival service offered by AWS. AWS S3 provides lifecycle rules to archive and restore objects from S3 to Glacier. Once the object is archived their storage class will change to Glacier. If the user sends a request for restore, the storage class will still be Glacier for the restored object. The user will be paying for both the archived copy as well as for the restored object. The object is available only for the duration specified in the restore request and if the user wants to modify that period, he has to raise another restore request with the updated duration. QUESTION 219 A user has created a public subnet with VPC and launched an EC2 instance within it. The user is trying to delete the subnet. What will happen in this scenario? A. It will delete the subnet and make the EC2 instance as a part of the default subnet B. It will not allow the user to delete the subnet until the instances are terminated C. It will delete the subnet as well as terminate the instances D. The subnet can never be deleted independently, but the user has to delete the VPC first Answer: B Explanation: A Virtual Private Cloud (VPC. is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. When an instance is launched it will have a network interface attached with it. The user cannot delete the subnet until he terminates the instance and deletes the network interface. QUESTION 220 A storage admin wants to encrypt all the objects stored in S3 using server side encryption. The user does not want to use the AES 256 encryption key provided by S3. How can the user achieve this? A. The admin should upload his secret key to the AWS console and let S3 decrypt the objects B. The admin should use CLI or API to upload the encryption key to the S3 bucket. When making a call to the S3 API mention the encryption key URL in each request C. S3 does not support client supplied encryption keys for server side encryption D. The admin should send the keys and encryption algorithm with each API call Answer: D 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. Amazon S3 never stores the user's encryption key. The user has to supply it for each encryption or decryption call. Latest AWS-SysOps questions and answers from Amazon Exam Center offered by Lead2pass for free share now! Read and remember all real questions answers, Guarantee pass AWS-SysOps real test 100% or full money back! AWS-SysOps new questions on Google Drive: <https://drive.google.com/open?id=0B3Syig5i8gpDekE1aUpSVGNHbWM> 2017 Amazon AWS-SysOps exam dumps (All 332 Q&As) from Lead2pass: <http://www.lead2pass.com/aws-sysops.html> [100% Exam Pass Guaranteed]