



Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt. ----- Guaranteed.



DOP-C01 Dumps
DOP-C01 Braindumps
DOP-C01 Real Questions
DOP-C01 Practice Test
DOP-C01 Actual Questions



Amazon

DOP-C01

AWS DevOps Engineer Professional (DOP-C01)



<https://killexams.com/pass4sure/exam-detail/DOP-C01>

Question #516

Which statement is true about configuring proxy support for Amazon Inspector agent on a Windows-based system?

- A. Amazon Inspector agent supports proxy usage on Windows-based systems through the use of the WinHTTP proxy.
- B. Amazon Inspector agent supports proxy usage on Linux-based systems but not on Windows.
- C. Amazon Inspector proxy support on Windows-based systems is achieved through installing proxy-enabled version of the agent which comes with preconfigured files that you need to edit to match your environment.
- D. Amazon Inspector agent supports proxy usage on Windows-based systems through awsagent.env configuration file.

Answer: A

Proxy support for AWS agents is achieved through the use of the WinHTTP proxy.

Reference:

https://docs.aws.amazon.com/inspector/latest/userguide/inspector_agents-on-win.html#inspectoragent-proxy

Question #517

What is the default maximum number of Roles per AWS account?

- A. 500
- B. 250
- C. 100
- D. There is no limit.

Answer: B

The default maximum number of Roles per AWS account is 250.

Reference:

<http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.htm>

Question #518

You have an application which consists of EC2 instances in an Auto Scaling group. Between a particular time frame every day, there is an increase in traffic to your website. Hence users are complaining of a poor response time on the application. You have configured your Auto Scaling group to deploy one new EC2 instance when CPU utilization is greater than 60% for 2 consecutive periods of 5 minutes.

What is the least cost-effective way to resolve this problem?

- A. Decrease the consecutive number of collection periods
- B. Increase the minimum number of instances in the Auto Scaling group
- C. Decrease the collection period to ten minutes
- D. Decrease the threshold CPU utilization percentage at which to deploy a new instance

Answer: B

If you increase the minimum number of instances, then they will be running even though the load is not high on the website. Hence you are incurring cost even though there is no need. All of the remaining options are possible options which can be used to increase the number of instances on a high load. For more information on On-demand scaling, please refer to the below link.

Reference:

<http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scale-based-on-demand.html>

Question #519

You have decided that you need to change the instance type of your production instances which are running as part of an AutoScaling group. The entire architecture is deployed using CloudFormation Template. You currently have 4 instances in Production. You cannot have any interruption in service and need to ensure 2 instances are always running during the update. Which of the options below listed can be used for this?

- A. AutoScalingRollingUpdate
- B. AutoScalingScheduledAction
- C. AutoScalingReplacingUpdate
- D. AutoScalingIntegrationUpdate

Answer: A

The AWS::AutoScaling::AutoScalingGroup resource supports an UpdatePolicy attribute. This is used to define how an Auto Scaling group resource is updated when an update to the Cloud Formation stack occurs. A common approach to updating an Auto Scaling group is to perform a rolling update, which is done by specifying the AutoScalingRollingUpdate policy. This retains the same Auto Scaling group and replaces old instances with new ones, according to the parameters specified. For more information on Autoscaling updates, please refer to the below link.

Reference:

<https://aws.amazon.com/premiumsupport/knowledge-center/auto-scaling-group-rolling-updates/>

Question #520

You currently have the following setup in AWS:

- 1) An Elastic Load Balancer
 - 2) Auto Scaling Group which launches EC2 Instances
 - 3) AMIs with your code pre-installed
- You want to deploy the updates of your app to only a certain number of users. You want to have a cost-effective solution. You should also be able to revert back quickly.

Which of the below solutions is the most feasible one?

- A. Create a second ELB, and a new Auto Scaling Group assigned a new Launch Configuration. Create a new AMI with the updated app. Use Route53 Weighted Round Robin records to adjust the proportion of traffic hitting the two ELBs.
- B. Create new AMIs with the new app. Then use the new EC2 instances in half proportion to the older instances.
- C. Redeploy with AWS Elastic Beanstalk and Elastic Beanstalk versions. Use Route 53 Weighted Round Robin records to adjust the proportion of traffic hitting the two ELBs
- D. Create a full second stack of instances, cut the DNS over to the new stack of instances, and change the DNS back if a rollback is needed.

Answer: A

The Weighted Routing policy of Route53 can be used to direct a proportion of traffic to your application. The best option is to create a second CLB, attach the new Autoscaling Group and then use Route53 to divert the traffic. Option B is wrong because just having EC2 instances running with the new code will not help. Option

C is wrong because Elastic beanstalk is good for development environments, and also there is no mention of having 2 environments where environment urls can be swapped. Option D is wrong because you still need Route53 to split the traffic.

Question #521

You have an application running a specific process that is critical to the application's functionality, and have added the health check process to your Auto Scaling

Group. The instances are showing healthy but the application itself is not working as it should. What could be the issue with the health check, since it is still showing the instances as healthy.

- A. You do not have the time range in the health check properly configured
- B. It is not possible for a health check to monitor a process that involves the application
- C. The health check is not configured properly

- D. The health check is not checking the application process

Answer: D

If you have custom health checks, you can send the information from your health checks to Auto Scaling so that Auto Scaling can use this information. For example, if you determine that an instance is not functioning as expected, you can set the health status of the instance to Unhealthy. The next time that Auto Scaling performs a health check on the instance, it will determine that the instance is unhealthy and then launch a replacement instance.

Question #522

You have just recently deployed an application on EC2 instances behind an ELB. After a couple of weeks, customers are complaining on receiving errors from the application. You want to diagnose the errors and are trying to get errors from the ELB access logs. But the ELB access logs are empty. What is the reason for this.

- A. You do not have the appropriate permissions to access the logs
- B. You do not have your CloudWatch metrics correctly configured
- C. ELB Access logs are only available for a maximum of one week
- D. Access logging is an optional feature of Elastic Load Balancing that is disabled by default

Answer: D

Elastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot issues. Access logging is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logging for your load balancer, Elastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify. You can disable access logging at any time.

Question #523

You have deployed an application to AWS which makes use of Autoscaling to launch new instances. You now want to change the instance type for the new instances. Which of the following is one of the action items to achieve this deployment?

- A. Use Elastic Beanstalk to deploy the new application with the new instance type
- B. Use CloudFormation to deploy the new application with the new instance type
- C. Create a new launch configuration with the new instance type
- D. Create new EC2 instances with the new instance type and attach it to the Autoscaling Group

Answer: C

The ideal way is to create a new launch configuration, attach it to the existing Auto Scaling group, and terminate the running instances. Option A is invalid because Elastic beanstalk cannot launch new instances on demand. Since the current scenario requires Autoscaling, this is not the ideal option. Option B is invalid because this will be a maintenance overhead, since you just have an Autoscaling Group. There is no need to create a whole CloudFormation template for this. Option D is invalid because Autoscaling Group will still launch EC2 instances with the older launch configuration.

Question #524

Your application stores sensitive information on an EBS volume attached to your EC2 instance. How can you protect your information? (Choose two.)

- A. Unmount the EBS volume, take a snapshot and encrypt the snapshot. Re-mount the Amazon EBS volume.
- B. It is not possible to encrypt an EBS volume, you must use a lifecycle policy to transfer data to S3 for encryption.
- C. Copy the unencrypted snapshot and check the box to encrypt the new snapshot. Volumes restored from this encrypted snapshot will also be encrypted.
- D. Create and mount a new, encrypted Amazon EBS volume. Move the data to the new volume. Delete the old Amazon EBS volume.

Answer: CD

These steps are given in the AWS documentation

To migrate data between encrypted and unencrypted volumes

- 1) Create your destination volume (encrypted or unencrypted, depending on your need).
- 2) Attach the destination volume to the instance that hosts the data to migrate.
- 3) Make the destination volume available by following the procedures in Making an Amazon EBS Volume Available for Use. For Linux instances, you can create a mount point at /mnt/destination and mount the destination volume there.
- 4) Copy the data from your source directory to the destination volume. It may be most convenient to use a bulk-copy utility for this.

To encrypt a volume's data by means of snapshot copying

- 1) Create a snapshot of your unencrypted CBS volume. This snapshot is also unencrypted.
- 2) Copy the snapshot while applying encryption parameters. The resulting target snapshot is encrypted.
- 3) Restore the encrypted snapshot to a new volume, which is also encrypted.

Question #525

Which Auto Scaling process would be helpful when testing new instances before sending traffic to them, while still keeping them in your Auto Scaling Group?

- A. Suspend the process AZ Rebalance
- B. Suspend the process Health Check
- C. Suspend the process Replace Unhealthy
- D. Suspend the process AddToLoadBalancer

Answer: D

If you suspend Add To Load Balancer, Auto Scaling launches the instances but does not add them to the load balancer or target group. If you resume the AddTo

Load Balancer process. Auto Scaling resumes adding instances to the load balancer or target group when they are launched. However, Auto Scaling does not add the instances that were launched while this process was suspended. You must register those instances manually. Option A is invalid because this just balances the number of CC2 instances in the group across the Availability Zones in the region Option B is invalid because this just checks the health of the instances. Auto

Scaling marks an instance as unhealthy if Amazon CC2 or Elastic Load Balancing tells Auto Scaling that the instance is unhealthy. Option C is invalid because this process just terminates instances that are marked as unhealthy and later creates new instances to replace them.

Question #526

You have an ELB setup in AWS with EC2 instances running behind it. You have been requested to monitor the incoming connections to the ELB.

Which of the below options can suffice this requirement?

- A. Use AWSCloudTrail with your load balancer
- B. Enable access logs on the load balancer
- C. Use a CloudWatch Logs Agent
- D. Create a custom metric CloudWatch filter on your load balancer

Answer: B

Elastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot issues.

Option A is invalid because this service will monitor all AWS services Option C and D are invalid since CLB already provides a logging feature.

Question #527

A DevOps Engineer has been asked to recommend a tool to deploy the components of a three-tier web application. This application will use Amazon DynamoDB as a database. Which deployment requires the LEAST amount of operational management?

- A. Use AWS CloudFormation to create a Classic Load Balancer and an Auto Scaling group. Use AWS OpsWorks to create the application and database resources. Deploy application updates with OpsWorks using lifecycle events.
- B. Use AWS OpsWorks to create a Classic Load Balancer, an Auto Scaling group application, and database resources. Deploy application updates using OpsWorks lifecycle events.
- C. Use AWS OpsWorks to create a Classic Load Balancer, Auto Scaling, and application resources. Use AWS CloudFormation to create the database resources. Deploy application updates using CloudFormation rolling updates.
- D. Use AWS CloudFormation to create a Classic Load Balancer, an Auto Scaling group, and database resources. Deploy application updates using CloudFormation rolling updates.

Answer: B

Question #528

A company uses AWS CodePipeline to manage and deploy infrastructure as code. The infrastructure is defined in AWS CloudFormation templates and is primarily comprised of multiple Amazon EC2 instances and Amazon RDS databases. The Security team has observed many operators creating inbound security group rules with a source CIDR of 0.0.0.0/0 and would like to proactively stop the deployment of rules with open CIDRs. The DevOps Engineer will implement a pre-deployment step that runs some security checks over the CloudFormation template before the pipeline processes it. This check should allow only inbound security group rules with a source CIDR of 0.0.0.0/0 if the rule has the description "Security Approval Ref XXXXX" (where XXXXX is a pre-allocated reference). The pipeline step should fail if this condition is not met and the deployment should be blocked. How should this be accomplished?

- A. Enable a SCP in AWS Organizations. The policy should deny access to the API call `CreateSecurityGroupRule` if the rule specifies 0.0.0.0/0 without a description referencing a security approval.
- B. Add an initial stage to CodePipeline called Security Check. This stage should call an AWS Lambda function that scans the CloudFormation template and fails the pipeline if it finds 0.0.0.0/0 in a security group without a description referencing a security approval.
- C. Create an AWS Config rule that is triggered on creation or edit of resource type `EC2::SecurityGroup`. This rule should call an AWS Lambda function to send a failure notification if the security group has any rules with a source CIDR of 0.0.0.0/0 without a description referencing a security approval.
- D. Modify the IAM role used by CodePipeline. The IAM policy should deny access.

Answer: B



SAMPLE QUESTIONS

*These questions are for demo purpose only. **Full version** is up to date and contains actual questions and answers.*

Killexams.com is an online platform that offers a wide range of services related to certification exam preparation. The platform provides actual questions, exam dumps, and practice tests to help individuals prepare for various certification exams with confidence. Here are some key features and services offered by Killexams.com:

Actual Exam Questions: *Killexams.com provides actual exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these actual questions, candidates can familiarize themselves with the content and format of the real exam.*

Exam Dumps: *Killexams.com offers exam dumps in PDF format. These dumps contain a comprehensive collection of questions and answers that cover the exam topics. By using these dumps, candidates can enhance their knowledge and improve their chances of success in the certification exam.*

Practice Tests: *Killexams.com provides practice tests through their desktop VCE exam simulator and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice tests cover a wide range of questions and enable candidates to identify their strengths and weaknesses.*

Guaranteed Success: *Killexams.com offers a success guarantee with their exam dumps. They claim that by using their materials, candidates will pass their exams on the first attempt or they will refund the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exams.*

Updated Content: *Killexams.com regularly updates its question bank and exam dumps to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.*

Technical Support: *Killexams.com provides free 24x7 technical support to assist candidates with any queries or issues they may encounter while using their services. Their certified experts are available to provide guidance and help candidates throughout their exam preparation journey.*

For More exams visit <https://killexams.com/vendors-exam-list>
Kill your exam at First Attempt....Guaranteed!