

Valid Test AWS-DevOps-Engineer-Professional Vce Free | AWS-DevOps-Engineer-Professional Test Study Guide



BTW, DOWNLOAD part of Test4Sure AWS-DevOps-Engineer-Professional dumps from Cloud Storage:
https://drive.google.com/open?id=1mrXy09mPQtOjX1Yv6_aT9unQmFnfyF3m

With “reliable credit” as the soul of our AWS-DevOps-Engineer-Professional study tool, “utmost service consciousness” as the management philosophy, we endeavor to provide customers with high quality service. Our customer service staff, who are willing to be your little helper and answer your any questions about our AWS Certified DevOps Engineer - Professional qualification test, fully implement the service principle of customer-oriented service activities, aiming at comprehensive, coordinated and sustainable cooperation relationship with every users. Any puzzle about our AWS-DevOps-Engineer-Professional Test Torrent will receive timely and effective response, just leave a message on our official website or send us an e-mail at your convenience.

Amazon DOP-C01 (AWS Certified DevOps Engineer - Professional) Exam is a professional-level certification that validates the candidate's expertise in DevOps practices and principles on the AWS platform. AWS Certified DevOps Engineer - Professional certification exam is designed for experienced DevOps engineers who have a deep understanding of the AWS services and can design, deploy, and manage highly scalable and fault-tolerant systems on the cloud. AWS-DevOps-Engineer-Professional Exam covers a range of topics, including continuous integration and delivery, infrastructure as code, monitoring and logging, security and compliance, and automation.

Amazon DOP-C01 (AWS Certified DevOps Engineer - Professional) Certification Exam is a professional-level exam offered by Amazon Web Services (AWS) that validates an individual's expertise in designing, deploying, and managing distributed systems and applications on the AWS platform. AWS Certified DevOps Engineer - Professional certification is designed for individuals who have already achieved the AWS Certified Developer - Associate or AWS Certified SysOps Administrator - Associate certifications and have at least two years of experience working with AWS.

>> Valid Test AWS-DevOps-Engineer-Professional Vce Free <<

AWS-DevOps-Engineer-Professional Test Study Guide | AWS-DevOps-Engineer-Professional Valid Exam Vce Free

Real Amazon AWS-DevOps-Engineer-Professional test questions provide the necessary knowledge and skills to clear the test in a

short time. When applicants don't prepare with the latest AWS Certified DevOps Engineer - Professional (AWS-DevOps-Engineer-Professional) exam questions they fail and lose money. Test4Sure provides valid AWS-DevOps-Engineer-Professional practice test material for applicants who want to pass the AWS-DevOps-Engineer-Professional exam quickly.

The AWS Certified DevOps Engineer - Professional certification is ideal for professionals who are experienced in AWS DevOps practices and want to advance their careers in this field. With this certification, professionals can demonstrate their expertise in designing and managing DevOps solutions using AWS technologies. AWS Certified DevOps Engineer - Professional certification is recognized globally and can open up new career opportunities for professionals in various industries.

Amazon AWS Certified DevOps Engineer - Professional Sample Questions (Q264-Q269):

NEW QUESTION # 264

A DevOps Engineer is leading the implementation for automating patching of Windows-based workstations in a hybrid cloud environment by using AWS Systems Manager (SSM).

What steps should the Engineer follow to set up Systems Manager to automate patching in this environment?

(Choose two.)

- A. Create an IAM service role for Systems Manager so that the `ssm.amazonaws.com` service can execute the `AssumeRole` operation. Register the role to enable the creation of a service token. Perform managed- instance activation with the newly created service role.
- B. Create multiple IAM service roles for Systems Manager so that the `ssm.amazonaws.com` service can execute the `AssumeRole` operation on every instance. Register the role on a per-resource level to enable the creation of a service token. Perform managed-instance activation with the newly created service role attached to each managed instance.
- C. Using previously obtained activation codes and activation IDs, download and install the SSM Agent on the hybrid servers, and register the servers or virtual machines on the Systems Manager service. Hybrid instances will show with an "i-" prefix in the SSM console as if they were provisioned as a regular Amazon EC2 instance.
- D. Run AWS Config to create a list of instances that are unpatched and not compliant. Create an instance scheduler job, and through an AWS Lambda function, perform the instance patching to bring them up to compliance.
- E. Using previously obtained activation codes and activation IDs, download and install the SSM Agent on the hybrid servers, and register the servers or virtual machines on the Systems Manager service. Hybrid instances will show with an "mi-" prefix in the SSM console.

Answer: A,D

NEW QUESTION # 265

Management has reported an increase in the monthly bill from Amazon Web Services, and they are extremely concerned with this increased cost. Management has asked you to determine the exact cause of this increase. After reviewing the billing report, you notice an increase in the data transfer cost. How can you provide management with a better insight into data transfer use?

- A. Using Amazon CloudWatch metrics, pull your Elastic Load Balancing outbound data transfer metrics monthly, and include them with your billing report to show which application is causing higher bandwidth usage.
- B. Update your Amazon CloudWatch metrics to use five-second granularity, which will give better detailed metrics that can be combined with your billing data to pinpoint anomalies.
- C. Deliver custom metrics to Amazon CloudWatch per application that breaks down application data transfer into multiple, more specific data points.
- D. Use Amazon CloudWatch Logs to run a map-reduce on your logs to determine high usage and data transfer.

Answer: C

Explanation:

You can publish your own metrics to CloudWatch using the AWS CLI or an API. You can view statistical graphs of your published metrics with the AWS Management Console. CloudWatch stores data about a metric as a series of data points. Each data point has an associated time stamp. You can even publish an aggregated set of data points called a statistic set. If you have custom metrics specific to your application, you can give a breakdown to the management on the exact issue.

Option A won't be sufficient to provide better insights. Option B is an overhead when you can make the application publish custom metrics. Option D is invalid because just the ELB metrics will not give the entire picture.

For more information on custom metrics, please refer to the below document link: from AWS

<http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.html>

NEW QUESTION # 266

A customer owns a simple API for their website that receives about 1,000 requests each day and has an average response time of 50 ms. It is currently hosted on one c4 large instance.

Which changes to the architecture will provide high availability at the LOWEST cost?

- A. Create an Auto Scaling group with a minimum of one instance and a maximum of two instances, then use an Application Load Balancer to balance the traffic.
- B. Recreate the API using Amazon API Gateway and use AWS Lambda as the service backend.
- C. Create an Auto Scaling group with a maximum of two instances, then use an Application Load Balancer to balance the traffic.
- D. Recreate the API using Amazon API Gateway and integrate the new API with the existing backend service.

Answer: A

NEW QUESTION # 267

You have an Autoscaling Group which is launching a set of t2.small instances. You now need to replace those instances with a larger instance type. How would you go about making this change in an ideal manner?

- A. Change the Instance type in the current launch configuration to the new instance type.
- B. Change the Instance type of the Underlying EC2 instance directly.
- C. Create another Autoscaling Group and attach the new instance type.
- D. Create a new launch configuration with the new instance type and update your Autoscaling Group.

Answer: D

Explanation:

Explanation

Answer - C

The AWS Documentation mentions

A launch configuration is a template that an Auto Scaling group uses to launch EC2 instances. When you create a launch configuration, you specify information for the instances such as the ID of the Amazon Machine Image (AMI), the instance type, a key pair, one or more security groups, and a block device mapping. If you've launched an EC2 instance before, you specified the same information in order to launch the instance.

When you create an Auto Scaling group, you must specify a launch configuration. You can specify your launch configuration with multiple Auto Scaling groups.

However, you can only specify one launch configuration for an Auto Scaling group at a time, and you can't modify a launch configuration after you've created it.

Therefore, if you want to change the launch configuration for your Auto Scaling group, you must create a launch configuration and then update your Auto Scaling group with the new launch configuration.

For more information on launch configurations please see the below link:

* http://docs.aws.amazon.com/autoscaling/latest/userguide/l_launchConfiguration.html

NEW QUESTION # 268

An ecommerce company is receiving reports that its order history page is experiencing delays in reflecting the processing status of orders. The order processing system consists of an AWS Lambda function using reserved concurrency. The Lambda function processes order messages from an Amazon SQS queue and inserts processed orders into an Amazon DynamoDB table. The DynamoDB table has Auto Scaling enabled for read and write capacity.

Which actions will diagnose and resolve the delay? (Select TWO.)

- A. Check the ApproximateAgeOfOldestMessage metric for the SQS queue and configure a redrive policy on the SQS queue.
- B. Check the Throttles metric for the Lambda function and increase the Lambda function timeout.
- C. Check the ApproximateAgeOfOldestMessage metric for the SQS queue and increase the Lambda function concurrency limit.
- D. Check the ThrottledWriteRequests metric for the DynamoDB table and increase the maximum write capacity units for the table's Auto Scaling policy.

- Answer: A,C**

• • • • •

[illegible]

P.S. Free & New AWS-DevOps-Engineer-Professional dumps are available on Google Drive shared by Test4Sure:

https://drive.google.com/open?id=1mrXy09mPQtOjX1Yv6_aT9unQmFmfyF3m