



**Amazon**

## **Exam Questions AWS-Certified-DevOps-Engineer-Professional**

Amazon AWS Certified DevOps Engineer Professional

#### NEW QUESTION 1

You run a clustered NoSQL database on AWS EC2 using AWS EBS. You need to reduce latency for database response times. Performance is the most important concern, not availability. You did not perform the initial setup, someone without much AWS knowledge did, so you are not sure if they configured everything optimally. Which of the following is NOT likely to be an issue contributing to increased latency?

- A. The EC2 instances are not EBS Optimized.
- B. The database and requesting system are both in the wrong Availability Zone.
- C. The EBS Volumes are not using PIOPS.
- D. The database is not running in a placement group

**Answer: B**

#### Explanation:

For the highest possible performance, all instances in a clustered database like this one should be in a single Availability Zone in a placement group, using EBS optimized instances, and using PIOPS SSD EBS Volumes. The particular Availability Zone the system is running in should not be important, as long as it is the same as the requesting resources.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

#### NEW QUESTION 2

Your system automatically provisions EIPs to EC2 instances in a VPC on boot. The system provisions the whole VPC and stack at once. You have two of them per VPC. On your new AWS account, your attempt to create a Development environment failed, after successfully creating Staging and Production environments in the same region. What happened?

- A. You didn't choose the Development version of the AMI you are using.
- B. You didn't set the Development flag to true when deploying EC2 instances.
- C. You hit the soft limit of 5 EIPs per region and requested a 6th.
- D. You hit the soft limit of 2 VPCs per region and requested a 3rd

**Answer: C**

#### Explanation:

There is a soft limit of 5 EIPs per Region for VPC on new accounts. The third environment could not allocate the 6th EIP.

Reference: [http://docs.aws.amazon.com/general/latest/gr/aws\\_service\\_limits.html#limits\\_vpc](http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_vpc)

#### NEW QUESTION 3

To monitor API calls against our AWS account by different users and entities, we can use to create a history of calls in bulk for later review, and use for reacting to AWS API calls in real-time.

- A. AWS Config; AWS Inspector
- B. AWS CloudTrail; AWS Config
- C. AWS CloudTrail; CloudWatch Events
- D. AWS Config; AWS Lambda

**Answer: C**

#### Explanation:

CloudTrail is a batch API call collection service, CloudWatch Events enables real-time monitoring of calls through the Rules object interface.

Reference: <https://aws.amazon.com/whitepapers/security-at-scale-governance-in-aws/>

#### NEW QUESTION 4

What is the scope of an EC2 security group?

- A. Availability Zone
- B. Placement Group
- C. Region
- D. VPC

**Answer: C**

#### Explanation:

A security group is tied to a region and can be assigned only to instances in the same region. You can't enable an instance to communicate with an instance outside its region using security group rules. Traffic from an instance in another region is seen as WAN bandwidth.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/resources.html>

#### NEW QUESTION 5

Which EBS volume type is best for high performance NoSQL cluster deployments?

- A. io1
- B. gp1
- C. standard
- D. gp2

**Answer: A**

#### Explanation:

io1 volumes, or Provisioned IOPS (PIOPS) SSDs, are best for: Critical business applications that require sustained IOPS performance, or more than 10,000 IOPS or 160 MiB/s of throughput per volume, like large database workloads, such as MongoDB.  
Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html>

#### NEW QUESTION 6

You are building out a layer in a software stack on AWS that needs to be able to scale out to react to increased demand as fast as possible. You are running the code on EC2 instances in an Auto Scaling Group behind an ELB. Which application code deployment method should you use?

- A. SSH into new instances that come online, and deploy new code onto the system by pulling it from an S3 bucket, which is populated by code that you refresh from source control on new pushes.
- B. Bake an AMI when deploying new versions of code, and use that AMI for the Auto Scaling Launch Configuration.
- C. Create a Dockerfile when preparing to deploy a new version to production and publish it to S3. Use UserData in the Auto Scaling Launch configuration to pull down the Dockerfile from S3 and run it when new instances launch.
- D. Create a new Auto Scaling Launch Configuration with UserData scripts configured to pull the latest code at all times.

**Answer: B**

#### Explanation:

the bootstrapping process can be slower if you have a complex application or multiple applications to install. Managing a fleet of applications with several build tools and dependencies can be a challenging task during rollouts. Furthermore, your deployment service should be designed to do faster rollouts to take advantage of Auto Scaling.

Reference: <https://d0.awsstatic.com/whitepapers/overview-of-deployment-options-on-aws.pdf>

#### NEW QUESTION 7

You need to perform ad-hoc analysis on log data, including searching quickly for specific error codes and reference numbers. Which should you evaluate first?

- A. AWS Elasticsearch Service
- B. AWS RedShift
- C. AWS EMR
- D. AWS DynamoDB

**Answer: A**

#### Explanation:

Amazon Elasticsearch Service (Amazon ES) is a managed service that makes it easy to deploy, operate, and scale Elasticsearch clusters in the AWS cloud. Elasticsearch is a popular open-source search and analytics engine for use cases such as log analytics, real-time application monitoring, and click stream analytics.

Reference:

<http://docs.aws.amazon.com/elasticsearch-service/latest/developerguide/what-is-amazon-elasticsearch-service.html>

#### NEW QUESTION 8

When thinking of AWS Elastic Beanstalk's model, which is true?

- A. Applications have many deployments, deployments have many environments.
- B. Environments have many applications, applications have many deployments.
- C. Applications have many environments, environments have many deployments.
- D. Deployments have many environments, environments have many application

**Answer: C**

#### Explanation:

Applications group logical services. Environments belong to Applications, and typically represent different deployment levels (dev, stage, prod, forth). Deployments belong to environments, and are pushes of bundles of code for the environments to run.

Reference: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/welcome.html>

#### NEW QUESTION 9

You are getting a lot of empty receive requests when using Amazon SQS. This is making a lot of unnecessary network load on your instances. What can you do to reduce this load?

- A. Subscribe your queue to an SNS topic instead.
- B. Use as long of a poll as possible, instead of short polls.
- C. Alter your visibility timeout to be shorter.
- D. Use `sqsd` on your EC2 instance

**Answer: B**

#### Explanation:

One benefit of long polling with Amazon SQS is the reduction of the number of empty responses, when there are no messages available to return, in reply to a ReceiveMessage request sent to an Amazon SQS queue. Long polling allows the Amazon SQS service to wait until a message is available in the queue before sending a response.

Reference:

<http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-long-polling.html>

#### NEW QUESTION 10

There are a number of ways to purchase compute capacity on AWS. Which orders the price per compute or memory unit from LOW to HIGH (cheapest to most expensive), on average?

- (A) On-Demand (B) Spot (C) Reserved

- A. A, B, C
- B. C, B, A
- C. B, C, A
- D. A, C, B

**Answer:** C

**Explanation:**

Spot instances are usually many, many times cheaper than on-demand prices. Reserved instances, depending on their term and utilization, can yield approximately 33% to 66% cost savings. On-Demand prices are the baseline price and are the most expensive way to purchase EC2 compute time. Reference: [https://d0.awsstatic.com/whitepapers/Cost\\_Optimization\\_with\\_AWS.pdf](https://d0.awsstatic.com/whitepapers/Cost_Optimization_with_AWS.pdf)

**NEW QUESTION 10**

You run operations for a company that processes digital wallet payments at a very high volume. One second of downtime, during which you drop payments or are otherwise unavailable, loses you on average USD 100. You balance the financials of the transaction system once per day. Which database setup is best suited to address this business risk?

- A. A multi-AZ RDS deployment with synchronous replication to multiple standbys and read-replicas for fast failover and ACID properties.
- B. A multi-region, multi-master, active-active RDS configuration using database-level ACID design principles with database trigger writes for replication.
- C. A multi-region, multi-master, active-active DynamoDB configuration using application control-level BASE design principles with change-stream write queue buffers for replication.
- D. A multi-AZ DynamoDB setup with changes streamed to S3 via AWS Kinesis, for highly durable storage and BASE properties.

**Answer:** C

**Explanation:**

Only the multi-master, multi-region DynamoDB answer makes sense. Multi-AZ deployments do not provide sufficient availability when a business loses USD 360,000 per hour of unavailability. As RDS does not natively support multi-region, and ACID does not perform well/at all over large distances between regions, only the DynamoDB answer works. Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.CrossRegionRepl.html>

**NEW QUESTION 12**

When thinking of DynamoDB, what are true of Local Secondary Key properties?

- A. Either the partition key or the sort key can be different from the table, but not both.
- B. Only the sort key can be different from the table.
- C. The partition key and sort key can be different from the table.
- D. Only the partition key can be different from the table.

**Answer:** B

**Explanation:**

Global secondary index — an index with a partition key and a sort key that can be different from those on the table. A global secondary index is considered "global" because queries on the index can span all of the data in a table, across all partitions. Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html>

**NEW QUESTION 16**

Which of these techniques enables the fastest possible rollback times in the event of a failed deployment?

- A. Rolling; Immutable
- B. Rolling; Mutable
- C. Canary or A/B
- D. Blue-Green

**Answer:** D

**Explanation:**

AWS specifically recommends Blue-Green for super-fast, zero-downtime deploys - and thus rollbacks, which are redeploying old code. You use various strategies to migrate the traffic from your current application stack (blue) to a new version of the application (green). This is a popular technique for deploying applications with zero downtime. Reference: <https://d0.awsstatic.com/whitepapers/overview-of-deployment-options-on-aws.pdf>

**NEW QUESTION 17**

Which of the following are not valid sources for OpsWorks custom cookbook repositories?

- A. HTTP(S)
- B. Git
- C. AWS EBS
- D. Subversion

**Answer:** C

**Explanation:**

Linux stacks can install custom cookbooks from any of the following repository types: HTTP or Amazon S3 archives. They can be either public or private, but Amazon S3 is typically the preferred option for a private archive. Git and Subversion repositories provide source control and the ability to have multiple versions. Reference: <http://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-installingcustom-enable.html>

#### NEW QUESTION 20

You are building a deployment system on AWS. You will deploy new code by bootstrapping instances in a private subnet in a VPC at runtime using UserData scripts pointing to an S3 zip file object, where your code is stored. An ELB in a public subnet has network interfaces and connectMty to the instances. Requests from users of the system are routed to the ELB via a Route53 A Record Alias. You do not use any VPC endpoints. Which is a risk of using this approach?

- A. Route53 Alias records do not always update dynamically with ELB network changes after deploys.
- B. If the NAT routing for the private subnet fails, deployments fail.
- C. Kernel changes to the base AMI may render the code inoperable.
- D. The instances cannot be in a private subnet if the ELB is in a public on

**Answer:** B

#### Explanation:

Since you are not using VPC endpoints, outbound requests for the code sitting in S3 are routed though the NAT for the VPC's private subnets. If this networking fails, runtime bootstrapping through code download will fail due to network unavailability and lack of access to the Internet, and thus Amazon S3. Reference: [http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC\\_NAT\\_Instance.html](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_NAT_Instance.html)

#### NEW QUESTION 24

Which major database needs a BYO license?

- A. PostgreSQL
- B. MariaDB
- C. MySQL
- D. Oracle

**Answer:** D

#### Explanation:

Oracle is not open source, and requires a bring your own license model. Reference: [http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Oracle.html](http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Oracle.html)

#### NEW QUESTION 27

What is the maximum supported single-volume throughput on EBS?

- A. 320MiB/s
- B. 160MiB/s
- C. 40MiB/s
- D. 640MiB/s

**Answer:** A

#### Explanation:

The ceiling throughput for PIOPS on EBS is 320MiB/s. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html>

#### NEW QUESTION 29

Your serverless architecture using AWS API Gateway, AWS Lambda, and AWS DynamoDB experienced a large increase in traffic to a sustained 400 requests per second, and dramatically increased in failure rates. Your requests, during normal operation, last 500 milliseconds on average. Your DynamoDB table did not exceed 50% of provisioned throughput, and Table primary keys are designed correctly. What is the most likely issue?

- A. Your API Gateway deployment is throttling your requests.
- B. Your AWS API Gateway Deployment is bottlenecking on request (de)serialization.
- C. You did not request a limit increase on concurrent Lambda function executions.
- D. You used Consistent Read requests on DynamoDB and are experiencing semaphore lock

**Answer:** C

#### Explanation:

AWS API Gateway by default throttles at 500 requests per second steady-state, and 1000 requests per second at spike. Lambda, by default, throttles at 100 concurrent requests for safety. At 500 milliseconds (half of a second) per request, you can expect to support 200 requests per second at 100 concurrency. This is less than the 400 requests per second your system now requires. Make a limit increase request via the AWS Support Console.

AWS Lambda: Concurrent requests safety throttle per account -> 100  
 Reference: [http://docs.aws.amazon.com/general/latest/gr/aws\\_service\\_limits.html#limits\\_lambda](http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_lambda)

#### NEW QUESTION 30

For AWS CloudFormation, which stack state refuses UpdateStack calls?

- A. `UPDATE_ROLLBACK_FAILED`
- B. `UPDATE_ROLLBACK_COMPLETE`
- C. `UPDATE_COMPLETE`
- D. `CREATE_COMPLETE`

**Answer:** A

#### Explanation:

When a stack is in the UPDATE\_ROLLBACK\_FAILED state, you can continue rolling it back to return it to a working state (to UPDATE\_ROLLBACK\_COMPLETE). You cannot update a stack that is in the UPDATE\_ROLLBACK\_FAILED state. However, if you can continue to roll it back, you can return the stack to its original

settings and try to update it again.

Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks-continueupdateorrollback.html>

#### NEW QUESTION 35

Which of these is not a Pseudo Parameter in AWS CloudFormation?

- A. AWS::StackName
- B. AWS::AccountId
- C. AWS::StackArn
- D. AWS::NotificationARNs

**Answer: C**

#### Explanation:

This is the complete list of Pseudo Parameters: AWS::AccountId, AWS::NotificationARNs, AWS::NoValue, AWS::Region, AWS::StackId, AWS::StackName

Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/pseudo-parameter-reference.html>

#### NEW QUESTION 39

Your API requires the ability to stay online during AWS regional failures. Your API does not store any state, it only aggregates data from other sources - you do not have a database. What is a simple but effective way to achieve this uptime goal?

- A. Use a CloudFront distribution to serve up your AP
- B. Even if the region your API is in goes down, the edge locations CloudFront uses will be fine.
- C. Use an ELB and a cross-zone ELB deployment to create redundancy across datacenter
- D. Even if a region fails, the other AZ will stay online.
- E. Create a Route53 Weighted Round Robin record, and if one region goes down, have that region redirect to the other region.
- F. Create a Route53 Latency Based Routing Record with Failover and point it to two identical deployments of your stateless API in two different region
- G. Make sure both regions use Auto Scaling Groups behind ELBs.

**Answer: D**

#### Explanation:

standard volumes, or Magnetic volumes, are best for: Cold workloads where data is infrequently accessed, or scenarios where the lowest storage cost is important.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html>

#### NEW QUESTION 43

When thinking of AWS OpsWorks, which of the following is not an instance type you can allocate in a stack layer?

- A. 24/7 instances
- B. Spot instances
- C. Time-based instances
- D. Load-based instances

**Answer: B**

#### Explanation:

AWS OpsWorks supports the following instance types, which are characterized by how they are started and stopped. 24/7 instances are started manually and run until you stop them. Time-based instances are run by AWS OpsWorks on a specified daily and weekly schedule. They allow your stack to automatically adjust the number of instances to accommodate predictable usage patterns. Load-based instances are automatically started and stopped by AWS OpsWorks, based on specified load metrics, such as CPU utilization. They allow your stack to automatically adjust the number of instances to accommodate variations in incoming traffic. Load-based instances are available only for Linux-based stacks. Reference: <http://docs.aws.amazon.com/opsworks/latest/userguide/welcome.html>

#### NEW QUESTION 46

You are building a mobile app for consumers to post cat pictures online. You will be storing the images in AWS S3. You want to run the system very cheaply and simply. Which one of these options allows you to build a photo sharing application without needing to worry about scaling expensive uploads processes, authentication/authorization and so forth?

- A. Build the application out using AWS Cognito and web identity federation to allow users to log in using Facebook or Google Account
- B. Once they are logged in, the secret token passed to that user is used to directly access resources on AWS, like AWS S3.
- C. Use JWT or SANIL compliant systems to build authorization policie
- D. Users log in with a username and password, and are given a token they can use indefinitely to make calls against the photo infrastructure.
- E. Use AWS API Gateway with a constantly rotating API Key to allow access from the client-sid
- F. Construct a custom build of the SDK and include S3 access in it.
- G. Create an AWS OAuth Service Domain and grant public signup and access to the domai
- H. During setup, add at least one major social media site as a trusted Identity Provider for users.

**Answer: A**

#### Explanation:

The short answer is that Amazon Cognito is a superset of the functionality provided by web identity federation. It supports the same providers, and you configure your app and authenticate with those providers in the same way. But Amazon Cognito includes a variety of additional features. For example, it enables your users to start using the app as a guest user and later sign in using one of the supported identity providers.

Reference:

<https://blogs.aws.amazon.com/security/post/Tx3SYCORF5EKRCO/How-Does-Amazon-Cognito-Relate-to-Existing-Web-Identity-Federatio>

#### NEW QUESTION 49

Your CTO has asked you to make sure that you know what all users of your AWS account are doing to change resources at all times. She wants a report of who is doing what over time, reported to her once per week, for as broad a resource type group as possible. How should you do this?

- A. Create a global AWS CloudTrail Trail
- B. Configure a script to aggregate the log data delivered to S3 once per week and deliver this to the CTO.
- C. Use CloudWatch Events Rules with an SNS topic subscribed to all AWS API call
- D. Subscribe the CTO to an email type delivery on this SNS Topic.
- E. Use AWS IAM credential reports to deliver a CSV of all uses of IAM User Tokens over time to the CTO.
- F. Use AWS Config with an SNS subscription on a Lambda, and insert these changes over time into a DynamoDB tabl
- G. Generate reports based on the contents of this table.

**Answer:** A

#### Explanation:

This is the ideal use case for AWS CloudTrail.

CloudTrail provides visibility into user activity by recording API calls made on your account. CloudTrail records important information about each API call, including the name of the API, the identity of the caller, the time of the API call, the request parameters, and the response elements returned by the AWS service. This information helps you to track changes made to your AWS resources and to troubleshoot operational issues. CloudTrail makes it easier to ensure compliance with internal policies and regulatory standards. Reference: <https://aws.amazon.com/Cloudtrail/faqs/>

#### NEW QUESTION 52

If I want CloudFormation stack status updates to show up in a continuous delivery system in as close to real time as possible, how should I achieve this?

- A. Use a long-poll on the Resources object in your CloudFormation stack and display those state changes in the UI for the system.
- B. Use a long-poll on the `ListStacks` API call for your CloudFormation stack and display those state changes in the UI for the system.
- C. Subscribe your continuous delivery system to an SNS topic that you also tell your CloudFormation stack to publish events into.
- D. Subscribe your continuous delivery system to an SQS queue that you also tell your CloudFormation stack to publish events into.

**Answer:** C

#### Explanation:

Use NotificationARNs.member.N when making a CreateStack call to push stack events into SNS in nearly real-time.

Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks-monitor-stack.html>

#### NEW QUESTION 55

There is a very serious outage at AWS. EC2 is not affected, but your EC2 instance deployment scripts stopped working in the region with the outage. What might be the issue?

- A. The AWS Console is down, so your CLI commands do not work.
- B. S3 is unavailable, so you can't create EBS volumes from a snapshot you use to deploy new volumes.
- C. AWS turns off the `DeployCode` API call when there are major outages, to protect from system floods.
- D. None of the other answers make sense
- E. If EC2 is not affected, it must be some other issue

**Answer:** B

#### Explanation:

S3 stores all snapshots. If S3 is unavailable, snapshots are unavailable.

Amazon EC2 also uses Amazon S3 to store snapshots (backup copies) of the data volumes. You can use snapshots for recovering data quickly and reliably in case of application or system failures. You can also use snapshots as a baseline to create multiple new data volumes, expand the size of an existing data volume, or move data volumes across multiple Availability Zones, thereby making your data usage highly scalable. For more information about using data volumes and snapshots, see Amazon Elastic Block Store.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonS3.html>

#### NEW QUESTION 58

You are designing a system which needs, at minimum, 8 m4.large instances operating to service traffic. When designing a system for high availability in the us-east-1 region, which has 6 Availability Zones, your company needs to be able to handle death of a full availability zone. How should you distribute the servers, to save as much cost as possible, assuming all of the EC2 nodes are properly linked to an ELB? Your VPC account can utilize us-east-1's AZ's a through f, inclusive.

- A. 3 servers in each of AZ's a through d, inclusive.
- B. 8 servers in each of AZ's a and b.
- C. 2 servers in each of AZ's a through e, inclusive.
- D. 4 servers in each of AZ's a through c, inclusive

**Answer:** C

#### Explanation:

You need to design for N+1 redundancy on Availability Zones.  $ZONE\_COUNT = (REQUIRED\_INSTANCES / INSTANCE\_COUNT\_PER\_ZONE) + 1$ . To minimize cost, spread the instances across as many possible zones as you can. By using a through e, you are allocating 5 zones. Using 2 instances, you have 10 total instances. If a single zone fails, you have 4 zones left, with 2 instances each, for a total of 8 instances. By spreading out as much as possible, you have increased cost by only 25% and significantly de-risked an availability zone failure.

Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html#concepts-regions-availability-zones>

#### NEW QUESTION 63

You have been asked to de-risk deployments at your company. Specifically, the CEO is concerned about outages that occur because of accidental inconsistencies

between Staging and Production, which sometimes cause unexpected behaviors in Production even when Staging tests pass.

You already use Docker to get high consistency between Staging and Production for the application environment on your EC2 instances. How do you further de-risk the rest of the execution environment, since in AWS, there are many service components you may use beyond EC2 virtual machines?

- A. Develop models of your entire cloud system in CloudFormation
- B. Use this model in Staging and Production to achieve greater parity.
- C. Use AWS Config to force the Staging and Production stacks to have configuration parity
- D. Any differences will be detected for you so you are aware of risks.
- E. Use AMIs to ensure the whole machine, including the kernel of the virtual machines, is consistent, since Docker uses Linux Container (LXC) technology, and we need to make sure the container environment is consistent.
- F. Use AWS ECS and Docker clustering
- G. This will make sure that the AMIs and machine sizes are the same across both environments.

**Answer:** A

**Explanation:**

Only CloudFormation's JSON Templates allow declarative version control of repeatably deployable models of entire AWS clouds.

Reference: <https://blogs.aws.amazon.com/application-management/blog/category/Best+practices>

**NEW QUESTION 65**

.....

## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### AWS-Certified-DevOps-Engineer-Professional Practice Exam Features:

- \* AWS-Certified-DevOps-Engineer-Professional Questions and Answers Updated Frequently
- \* AWS-Certified-DevOps-Engineer-Professional Practice Questions Verified by Expert Senior Certified Staff
- \* AWS-Certified-DevOps-Engineer-Professional Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* AWS-Certified-DevOps-Engineer-Professional Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
[Order The AWS-Certified-DevOps-Engineer-Professional Practice Test Here](#)