www.CertBus.com

CERTBUS

DOP-C01^{Q&As}

AWS Certified DevOps Engineer - Professional (DOP-C01)

Pass Amazon DOP-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.certbus.com/aws-devops-engineer-professional.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon Official Exam Center

⚙ **Instant Download** After Purchase

⚙ **100% Money Back** Guarantee

⚙ **365 Days** Free Update

⚙ **800,000+** Satisfied Customers

QUESTION 1

A DevOps Engineer has several legacy applications that all generate different log formats. The Engineer must standardize the formats before writing them to Amazon S3 for querying and analysis.

How can this requirement be met at the LOWEST cost?

A. Have the application send its logs to an Amazon EMR cluster and normalize the logs before sending them to Amazon S3

B. Have the application send its logs to Amazon QuickSight, then use the Amazon QuickSight SPICE engine to normalize the logs. Do the analysis directly from Amazon QuickSight

C. Keep the logs in Amazon S3 and use Amazon Redshift Spectrum to normalize the logs in place

D. Use Amazon Kinesis Agent on each server to upload the logs and have Amazon Kinesis Data Firehose use an AWS Lambda function to normalize the logs before writing them to Amazon S3

Correct Answer: D

QUESTION 2

A DevOps Engineer is designing a deployment strategy for a web application. The application will use an Auto Scaling group to launch Amazon EC2 instances using an AMI. The same infrastructure will be deployed in multiple environments (development, test, and quality assurance). The deployment strategy should meet the following requirements:

1.

 Minimize the startup time for the instance

2.

 Allow the same AMI to work in multiple environments

3.

 Store secrets for multiple environments securely

How should this be accomplished?

A. Preconfigure the AMI using an AWS Lambda function that launches an Amazon EC2 instance, and then runs a script to install the software and create the AMI. Configure an Auto Scaling lifecycle hook to determine which environment the instance is launched in, and, based on that finding, run a configuration script. Save the secrets on an .ini file and store them in Amazon S3. Retrieve the secrets using a configuration script in EC2 user data.

B. Preconfigure the AMI by installing all the software using AWS Systems Manager automation and configure Auto Scaling to tag the instances at launch with their specific environment. Then use a bootstrap script in user data to read the tags and configure settings for the environment. Use the AWS Systems Manager Parameter Store to store the secrets using AWS KMS.

C. Use a standard AMI from the AWS Marketplace. Configure Auto Scaling to detect the current environment. Install the software using a script in Amazon EC2 user data. Use AWS Secrets Manager to store the credentials for all

environments.

D. Preconfigure the AMI by installing all the software and configuration for all environments. Configure Auto Scaling to tag the instances at launch with their environment. Use the Amazon EC2 user data to trigger an AWS Lambda function that reads the instance ID and then reconfigures the setting for the proper environment. Use the AWS Systems Manager Parameter Store to store the secrets using AWS KMS.

Correct Answer: B

QUESTION 3

Which of the following is an invalid variable name in Ansible?

A. host1st_ref

B. host-first-ref

C. Host1stRef

D. host_first_ref

Correct Answer: B

Variable names can contain letters, numbers and underscores and should always start with a letter. Invalid variable examples, `host first ref\\', `1st_host_ref\\'\\'.

Reference: http://docs.ansible.com/ansible/playbooks_variables.html#what-makes-a-valid-variable-name

QUESTION 4

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

Correct Answer: C

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/

QUESTION 5

For AWS CloudFormation, which stack state refuses UpdateStack calls?

A. UPDATE_ROLLBACK_FAILED

B. UPDATE_ROLLBACK_COMPLETE

C. UPDATE_COMPLETE

D. CREATE_COMPLETE

Correct Answer: A

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- stackscontinueu pdaterollback.html

---

QUESTION 6

A production account has a requirement that any Amazon EC2 instance that has been logged into manually must be terminated within 24 hours. All applications in the production account are using Auto Scaling groups with Amazon

CloudWatch Logs agent configured.

How can this process be automated?

A. Create a CloudWatch Logs subscription to an AWS Step Functions application. Configure the function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned. Then create a CloudWatch Events rule to trigger a second AWS Lambda function once a day that will terminate all instances with this tag.

B. Create a CloudWatch alarm that will trigger on the login event. Send the notification to an Amazon SNS topic that the Operations team is subscribed to, and have them terminate the EC2 instance within 24 hours.

C. Create a CloudWatch alarm that will trigger on the login event. Configure the alarm to send to an Amazon SQS queue. Use a group of worker instances to process messages from the queue, which then schedules the Amazon CloudWatch Events rule to trigger.

D. Create a CloudWatch Logs subscription in an AWS Lambda function. Configure the function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned. Create a CloudWatch Events rule to trigger a daily Lambda function that terminates all instances with this tag.

Correct Answer: D

QUESTION 7

Your company needs to automate 3 layers of a large cloud deployment. You want to be able to track this deployment\\'s evolution as it changes over time, and carefully control any alterations. What is a good way to automate a stack to meet these requirements?

A. Use OpsWorks Stacks with three layers to model the layering in your stack.

B. Use CloudFormation Nested Stack Templates, with three child stacks to represent the three logical layers of your cloud.

C. Use AWS Config to declare a configuration set that AWS should roll out to your cloud.

D. Use Elastic Beanstalk Linked Applications, passing the important DNS entires between layers using the metadata interface.

Correct Answer: B

Only CloudFormation allows source controlled, declarative templates as the basis for stack automation. Nested Stacks help achieve clean separation of layers while simultaneously providing a method to control all layers at once when

needed.

Reference:

https://blogs.aws.amazon.com/application-management/post/Tx1T9JYQOS8AB9I/Use-Nested-Stacks-to- Create-Reusable-Templates-and-Support-Role-Specialization

QUESTION 8

A company has a single Developer writing code for an automated deployment pipeline. The Developer is storing source code in an Amazon S3 bucket for each project. The company wants to add more Developers to the team but is concerned about code conflicts and lost work. The company also wants to build a test environment to deploy newer versions of code for testing and allow Developers to automatically deploy to both environments when code is changed in the repository.

What is the MOST efficient way to meet these requirements?

A. Create an AWS CodeCommit repository for each project, use the master branch for production code, and create a testing branch for code deployed to testing. Use feature branches to develop new features and pull requests to merge code to testing and master branches.

B. Create another S3 bucket for each project for testing code, and use an AWS Lambda function to promote code changes between testing and production buckets. Enable versioning on all buckets to prevent code conflicts.

C. Create an AWS CodeCommit repository for each project, and use the master branch for production and test code with different deployment pipelines for each environment. Use feature branches to develop new features.

D. Enable versioning and branching on each S3 bucket, use the master branch for production code, and create a testing branch for code deployed to testing. Have Developers use each branch for developing in each environment.

Correct Answer: A

QUESTION 9

A company is using several AWS CloudFormation templates for deploying infrastructure as code. In most of the deployments, the company uses Amazon EC2 Auto Scaling groups. A DevOps Engineer needs to update the AMIs for the Auto

Scaling group in the template if newer AMIs are available.

How can these requirements be met?

A. Manage the AMI mappings in the CloudFormation template. Use Amazon CloudWatch Events for detecting new

AMIs and updating the mapping in the template. Reference the map in the launch configuration resource block.

B. Use conditions in the AWS CloudFormation template to check if new AMIs are available and return the AMI ID. Reference the returned AMI ID in the launch configuration resource block.

C. Use an AWS Lambda-backed custom resource in the template to fetch the AMI IDs. Reference the returned AMI ID in the launch configuration resource block.

D. Launch an Amazon EC2 m4.small instance and run a script on it to check for new AMIs. If new AMIs are available, the script should update the launch configuration resource block with the new AMI ID.

Correct Answer: D

Reference:

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/walkthrough-customresources- lambda-lookup-amiids.html

---

QUESTION 10

A company has 100 GB of log data in an Amazon S3 bucket stored in .csv format. SQL developers want to query this data and generate graphs to visualize it. They also need an efficient, automated way to store metadata from the .csv file.

Which combination of steps should be taken to meet these requirements with the LEAST amount of effort? (Choose three.)

A. Filter the data through AWS X-Ray to visualize the data.

B. Filter the data through Amazon QuickSight to visualize the data.

C. Query the data with Amazon Athena.

D. Query the data with Amazon Redshift.

E. Use AWS Glue as the persistent metadata store.

F. Use Amazon S3 as the persistent metadata store.

Correct Answer: BCF

QUESTION 11

How does Amazon RDS multi Availability Zone model work?

A. A second, standby database is deployed and maintained in a different availability zone from master, using synchronous replication.

B. A second, standby database is deployed and maintained in a different availability zone from master using asynchronous replication.

C. A second, standby database is deployed and maintained in a different region from master using asynchronous replication.

D. A second, standby database is deployed and maintained in a different region from master using synchronous replication.

Correct Answer: A

In a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone. Reference:
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html

QUESTION 12

A company wants to use AWS CloudFormation for infrastructure deployment. The company has strict tagging and resource requirements and wants to limit the deployment to two Regions. Developers will need to deploy multiple versions of the same application.

Which solution ensures resources are deployed in accordance with company policy?

A. Create AWS Trusted Advisor checks to find and remediate unapproved CloudFormation StackSets.

B. Create a CloudFormation drift detection operation to find and remediate unapproved CloudFormation StackSets.

C. Create CloudFormation StackSets with approved CloudFormation templates.

D. Create AWS Service Catalog products with approved CloudFormation templates.

Correct Answer: C

Reference: https://aws.amazon.com/blogs/aws/use-cloudformation-stacksets-to-provision-resources- across-multiple-aws-accounts-and-regions/

QUESTION 13

A DevOps Engineer is architecting a continuous development strategy for a company\\'s software as a service (SaaS) web application running on AWS. For application and security reasons, users subscribing to this application are distributed across multiple Application Load Balancers (ALBs), each of which has a dedicated Auto Scaling group and fleet of Amazon EC2 instances. The application does not require a build stage, and when it is committed to AWS CodeCommit, the application must trigger a simultaneous deployment to all ALBs, Auto Scaling groups, and EC2 fleets.

Which architecture will meet these requirements with the LEAST amount of configuration?

A. Create a single AWS CodePipeline pipeline that deploys the application in parallel using unique AWS CodeDeploy applications and deployment groups created for each ALB-Auto Scaling group pair.

B. Create a single AWS CodePipeline pipeline that deploys the application using a single AWS CodeDeploy application and single deployment group.

C. Create a single AWS CodePipeline pipeline that deploys the application in parallel using a single AWS CodeDeploy application and unique deployment group for each ALB-Auto Scaling group pair.

D. Create an AWS CodePipeline pipeline for each ALB-Auto Scaling group pair that deploys the application using an AWS CodeDeploy application and deployment group created for the same ALB- Auto Scaling group pair.

Correct Answer: C

QUESTION 14

A DevOps engineer is tasked with migrating Docker containers used for a workload to AWS. The solution must allow for changes to be deployed into development and test environments automatically by updating each container and checking it into a container registry. Once the containers are pushed, they must be deployed automatically.

Which solution will meet these requirements?

A. Store container images in Amazon S3. Run the containers in AWS Elastic Beanstalk using a multicontainer Docker environment. Configure Elastic Beanstalk to redeploy the containers if it detects a new version in Amazon S3.

B. Store container images in AWS Artifact. Use AWS CodePipeline to trigger a deployment if a new container version is created. Use AWS CodeDeploy to deploy new containers to Amazon EKS.

C. Store container images in Amazon ECR. Use AWS CodePipeline to trigger a deployment if a new container version is created. Use AWS CodeDeploy to deploy the image to AWS Fargate.

D. Store container images in Docker Hub. Install Docker on an Amazon EC2 instance and use AWS CodePipeline and AWS CodeDeploy to deploy any new containers.

Correct Answer: C

QUESTION 15

You run a SIP-based telephony application that uses Amazon EC2 for its web tier and uses MySQL on Amazon RDS as its database. The application stores only the authentication profile data for its existing users in the database and

therefore is read-intensive. Your monitoring system shows that your web instances and the database have high CPU utilization.

Which of the following steps should you take in order to ensure the continual availability of your application? (Choose two.)

A. Use a CloudFront RTMP download distribution with the application tier as the origin for the distribution.

B. Set up an Auto Scaling group for the application tier and a policy that scales based on the Amazon EC2 CloudWatch CPU utilization metric.

C. Vertically scale up the Amazon EC2 instances manually.

D. Set up an Auto Scaling group for the application tier and a policy that scales based on the Amazon RDS CloudWatch CPU utilization metric.

E. Switch to General Purpose (SSD) Storage from Provisioned IOPS Storage (PIOPS) for the Amazon RDS database.

F. Use multiple Amazon RDS read replicas.

Correct Answer: BF

[DOP-C01 VCE Dumps](#)   [DOP-C01 Practice Test](#)   [DOP-C01 Study Guide](#)