



Free Questions for DOP-C01 by vceexamstest

Shared by Newton on 18-01-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

An education company has a Docker-based application running on multiple Amazon EC2 instances in an Amazon ECS cluster. When deploying a new version of the application, the Developer, pushes a new image to a private Docker container registry, and then stops and starts all tasks to ensure that they all have the latest version of the application. The Developer discovers that the new tasks are occasionally running with an old image. How can this issue be prevented?

Options:

- A) After pushing the new image, restart ECS Agent, and then start the tasks.
- B) Use 'latest' for the Docker image tag in the task definition.
- C) Update the digest on the task definition when pushing the new image.
- D) Use Amazon ECR for a Docker container registry.

Answer:

C

Question 2

Question Type: MultipleChoice

A company must ensure consistent behavior of an application running on Amazon Linux in its corporate ecosystem before moving into AWS. The company has an existing automated server build system using VMware. The goal is to demonstrate the functionality of the application and its prerequisites on the new target operating system. The DevOps Engineer needs to use the existing corporate server pipeline and virtualization software to create a server image. The server image will be tested on- premises to resemble the build on Amazon EC2 as closely as possible. How can this be accomplished?

Options:

- A)** Download and integrate the latest ISO of CentOS 7 and execute the application deployment on the resulting server.
- B)** Launch an Amazon Linux AMI using an AWS OpsWorks deployment agent onto the on-premises infrastructure, then execute the application deployment.
- C)** Build an EC2 instance with the latest Amazon Linux operating system, and use the AWS Import/Export service to export the EC2 image to a VMware ISO in Amazon S3. Then import the resulting ISO onto the on-premises system.
- D)** Download and integrate the latest ISO of Amazon Linux 2 and execute the application deployment on the resulting server. Confirm that operating system testing results are consistent with EC2 operating system behavior.

Answer:

D

Question 3

Question Type: MultipleChoice

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? (Select TWO.)

Options:

- A)** 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.
- B)** 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.
- 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 'mi-' prefix in the SSM console.
- D)** 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.
- E)** 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.

Answer:

B, C

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/sysman-managed-instance-activation.html>

<https://docs.aws.amazon.com/systems-manager/latest/userguide/sysman-install-managed-win.html>

Question 4

Question Type: MultipleChoice

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: ' Minimize the startup time for the instance ' Allow the same AMI to work in multiple environments ' Store secrets for multiple environments securely How should this be accomplished?

Options:

- 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.

Answer:

A

Question 5

Question Type: MultipleChoice

The Development team at an online retailer has moved to Business support and want to take advantage of the AWS Health Dashboard and the AWS Health API to automate remediation actions for issues with the health of AWS resources. The first use case is to respond to AWS detecting an IAM access key that is listed on a public code repository site. The automated response will be to delete the IAM access key and send a notification to the Security team. How should this be achieved?

Options:

- A)** Create an AWS Lambda function to delete the IAM access key. Send AWS CloudTrail logs to AWS CloudWatch logs. Create a CloudWatch Logs metric filter for the `AWS_RISK_CREDENTIALS_EXPOSED` event with two actions: first, run the Lambda function; second, use Amazon SNS to send a notification to the Security team.
- B)** Create an AWS Lambda function to delete the IAM access key. Create an AWS Config rule for changes to `aws.health` and the `AWS_RISK_CREDENTIALS_EXPOSED` event with two actions: first, run the Lambda function; second, use Amazon SNS to send a notification to the Security team.
- C)** Use AWS Step Functions to create a function to delete the IAM access key, and then use Amazon SNS to send a notification to the Security team. Create an AWS Personal Health Dashboard rule for the `AWS_RISK_CREDENTIALS_EXPOSED` event; set the target of the Personal Health Dashboard rule to Step Functions.
- D)** Use AWS Step Functions to create a function to delete the IAM access key, and then use Amazon SNS to send a notification to the Security team. Create an Amazon CloudWatch Events rule with an `aws.health` event source and the `AWS_RISK_CREDENTIALS_EXPOSED` event, set the target of the CloudWatch Events rule to Step Functions.

Answer:

A

Question 6

Question Type: MultipleChoice

A company is migrating an application to AWS that runs on a single Amazon EC2 instance. Because of licensing limitations, the application does not support horizontal scaling. The application will be using Amazon Aurora for its database. How can the DevOps Engineer architect automated healing to automatically recover from EC2 and Aurora failures, in addition to recovering across Availability Zones (AZs), in the MOST cost-effective manner?

Options:

- A)** Create an EC2 Auto Scaling group with a minimum and maximum instance count of 1, and have it span across AZs. Use a single-node Aurora instance.
- B)** Create an EC2 instance and enable instance recovery. Create an Aurora database with a read replica in a second AZ, and promote it to a primary database instance if the primary database instance fails.
- C)** Create an Amazon CloudWatch Events rule to trigger an AWS Lambda function to start a new EC2 instance in an available AZ when the instance status reaches a failure state. Create an Aurora database with a read replica in a second AZ, and promote it to a primary database instance when the primary database instance fails.
- D)** Assign an Elastic IP address on the instance. Create a second EC2 instance in a second AZ. Create an Amazon CloudWatch Events rule to trigger an AWS Lambda function to move the Elastic IP address to the second instance when the first instance fails. Use a single-node Aurora instance.

Answer:

C

To Get Premium Files for DOP-C01 Visit

<https://www.p2pexams.com/products/dop-c01>

For More Free Questions Visit

<https://www.p2pexams.com/amazon/pdf/dop-c01>

