

Exam Questions DOP-C01

AWS Certified DevOps Engineer- Professional

<https://www.2passeasy.com/dumps/DOP-C01/>



NEW QUESTION 1

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

Explanation:

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. For more information on Autoscaling health checks, please refer to the below document link: from AWS <http://docs.aws.amazon.com/autoscaling/latest/userguide/healthcheck.html>

NEW QUESTION 2

Your company has multiple applications running on AWS. Your company wants to develop a tool that notifies on-call teams immediately via email when an alarm is triggered in your environment. You have multiple on-call teams that work different shifts, and the tool should handle notifying the correct teams at the correct times. How should you implement this solution?

- A. Create an Amazon SNS topic and an Amazon SQS queue
- B. Configure the Amazon SQS queue as a subscriber to the Amazon SNS topic. Configure CloudWatch alarms to notify this topic when an alarm is triggered
- C. Create an Amazon EC2 Auto Scaling group with both minimum and desired Instances configured to 0. Worker nodes in this group spawn when messages are added to the queue
- D. Workers then use Amazon Simple Email Service to send messages to your on-call teams.
- E. Create an Amazon SNS topic and configure your on-call team email addresses as subscriber
- F. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to this new topic
- G. Notifications will be sent to on-call users when a CloudWatch alarm is triggered.
- H. Create an Amazon SNS topic and configure your on-call team email addresses as subscriber
- I. Create a secondary Amazon SNS topic for alarms and configure your CloudWatch alarms to notify this topic when triggered
- J. Create an HTTP subscriber to this topic that notifies your application via HTTP POST when an alarm is triggered
- K. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to the first topic so that on-call engineers receive alerts.
- L. Create an Amazon SNS topic for each on-call group, and configure each of these with the team member emails as subscriber
- M. Create another Amazon SNS topic and configure your CloudWatch alarms to notify this topic when triggered
- N. Create an HTTP subscriber to this topic that notifies your application via HTTP POST when an alarm is triggered
- O. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to the correct team topic when on shift.

Answer: D

Explanation:

Option D fulfills all the requirements

1) First is to create a SNS topic for each group so that the required members get the email addresses.
2) Ensure the application uses the HTTPS endpoint and the SDK to publish messages. Option A is invalid because the SQS service is not required.
Option B and C are incorrect. As per the requirement we need to provide notification to only those on-call teams who are working in that particular shift when an alarm is triggered. It need not have to be sent to all the on-call teams of the company. With Option B & C, since we are not configuring the SNS topic for each on-call team the notifications will be sent to all the on-call teams. Hence these 2 options are invalid. For more information on setting up notifications, please refer to the below document link: from AWS http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US_SetupSNS.html

NEW QUESTION 3

You have an Auto Scaling group of Instances that processes messages from an Amazon Simple Queue Service (SQS) queue. The group scales on the size of the queue. Processing involves calling a third-party web service. The web service is complaining about the number of failed and repeated calls it is receiving from you. You have noticed that when the group scales in, instances are being terminated while they are processing. What cost-effective solution can you use to reduce the number of incomplete process attempts?

- A. Create a new Auto Scaling group with minimum and maximum of 2 and instances running web proxy software
- B. Configure the VPC route table to route HTTP traffic to these web proxies.
- C. Modify the application running on the instances to enable termination protection while it processes a task and disable it when the processing is complete.
- D. Increase the minimum and maximum size for the Auto Scaling group, and change the scaling policies so they scale less dynamically.
- E. Modify the application running on the instances to put itself into an Auto Scaling Standby state while it processes a task and return itself to InService when the processing is complete.

Answer: D

Explanation:

The following diagram shows the lifecycle of the instances in Autoscaling

You can put the instances in a standby state, via the application, do the processing and then put the instance back in a state where it can be governed by the Autoscaling Group.

For more information on the Autoscaling Group Lifecycle please refer to the below link:

<http://docs.aws.amazon.com/autoscaling/latest/userguide/AutoScalingGroupLifecycle.htm> | Note: As per AWS documentation.

To control whether an Auto Scaling group can terminate a particular instance when scaling in, use instance protection.

It is termed as Instance protection rather than termination protection when we refer it with "Scaling in process" of ASG.

For more information please view the following link: <https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-instance-termination.html> | Instance protection-

instance

NEW QUESTION 4

You have been requested to use CloudFormation to maintain version control and achieve automation for the applications in your organization. How can you best use CloudFormation to keep everything agile and maintain multiple environments while keeping cost down?

- A. Create separate templates based on functionality, create nested stacks with CloudFormation.
- B. Use CloudFormation custom resources to handle dependencies between stacks
- C. Create multiple templates in one CloudFormation stack.
- D. Combine all resources into one template for version control and automation.

Answer: A

Explanation:

As your infrastructure grows, common patterns can emerge in which you declare the same components in each of your templates. You can separate out these common components and create dedicated templates for them. That way, you can mix and match different templates but use nested stacks to create a single, unified stack. Nested stacks are stacks that create other stacks. To create nested stacks, use the AWS::CloudFormation::Stack resource in your template to reference

other templates. For more information on CloudFormation best practices please refer to the below link:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html>

NEW QUESTION 5

You are using Elastic Beanstalk to manage your e-commerce store. The store is based on an open source e-commerce platform and is deployed across multiple instances in an Auto Scaling group. Your development team often creates new "extensions" for the e-commerce store. These extensions include PHP source code as well as an SQL upgrade script used to make any necessary updates to the database schema. You have noticed that some extension deployments fail due to an error when running the SQL upgrade script. After further investigation, you realize that this is because the SQL script is being executed on all of your Amazon EC2 instances. How would you ensure that the SQL script is only executed once per deployment regardless of how many Amazon EC2 instances are running at the time?

- A. Use a "Container command" within an Elastic Beanstalk configuration file to execute the script, ensuring that the "leader only" flag is set to true.
- B. Make use of the Amazon EC2 metadata service to query whether the instance is marked as the leader in the Auto Scaling group
- C. Only execute the script if "true" is returned.
- D. Use a "Solo Command" within an Elastic Beanstalk configuration file to execute the script
- E. The Elastic Beanstalk service will ensure that the command is only executed once.
- F. Update the Amazon RDS security group to only allow write access from a single instance in the Auto Scaling group; that way, only one instance will successfully execute the script on the database.

Answer: A

Explanation:

You can use the container_commands key to execute commands that affect your application source code. Container commands run after the application and web server have been set up and the application version archive has been extracted, but before the application version is deployed. Non-container commands and other customization operations are performed prior to the application source code being extracted.

You can use leader_only to only run the command on a single instance, or configure a test to only run the command when a test command evaluates to true.

Leader-only container commands are only executed during environment creation and deployments, while other commands and server customization operations are performed every time an instance is provisioned or updated. Leader-only container commands are not executed due to launch configuration changes, such as a change in the AMI ID or instance type. For more information on customizing containers, please visit the below URL:

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/customize-containers-ec2.html>

NEW QUESTION 6

You have a current CloudFormation template defined in AWS. You need to change the current alarm threshold defined in the CloudWatch alarm. How can you achieve this?

- A. Currently, there is no option to change what is already defined in CloudFormation templates.
- B. Update the template and then update the stack with the new template
- C. Automatically all resources will be changed in the stack.
- D. Update the template and then update the stack with the new template
- E. Only those resources that need to be changed will be changed
- F. All other resources which do not need to be changed will remain as they are.
- G. Delete the current CloudFormation template
- H. Create a new one which will update the current resources.

Answer: C

Explanation:

Option A is incorrect because CloudFormation templates have the option to update resources.

Option B is incorrect because only those resources that need to be changed as part of the stack update are actually updated.

Option D is incorrect because deleting the stack is not the ideal option when you already have a change option available.

When you need to make changes to a stack's settings or change its resources, you update the stack instead of deleting it and creating a new stack. For example, if you

have a stack with an EC2 instance, you can update the stack to change the instance's AMI ID.

When you update a stack, you submit changes, such as new input parameter values or an updated template. AWS CloudFormation compares the changes you submit with the current state of your stack and updates only the changed resources

For more information on stack updates please refer to the below link:

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

NEW QUESTION 7

When an Auto Scaling group is running in Amazon Elastic Compute Cloud (EC2), your application rapidly scales up and down in response to load within a 10-minute window; however, after the load peaks, you begin to see problems in your configuration management system where previously terminated Amazon EC2

resources are still showing as active. What would be a reliable and efficient way to handle the cleanup of Amazon EC2 resources within your configuration management system? Choose two answers from the options given below

- A. Write a script that is run by a daily cron job on an Amazon EC2 instance and that executes API Describe calls of the EC2 Auto Scaling group and removes terminated instances from the configuration management system.
- B. Configure an Amazon Simple Queue Service (SQS) queue for Auto Scaling actions that has a script that listens for new messages and removes terminated instances from the configuration management system.
- C. Use your existing configuration management system to control the launching and bootstrapping of instances to reduce the number of moving parts in the automation.
- D. Write a small script that is run during Amazon EC2 instance shutdown to de-register the resource from the configuration management system.

Answer: AD

Explanation:

There is a rich brand of CLI commands available for EC2 Instances. The CLI is located in the following link:

- <http://docs.aws.amazon.com/cli/latest/reference/ec2/>

You can then use the describe instances command to describe the EC2 instances.

If you specify one or more instance IDs, Amazon EC2 returns information for those instances. If you do not specify instance IDs, Amazon EC2 returns information for all relevant instances. If you specify an instance ID that is not valid, an error is returned. If you specify an instance that you do not own, it is not included in the returned results.

- <http://docs.aws.amazon.com/cli/latest/reference/ec2/describe-instances.html>

You can use the EC2 instances to get those instances which need to be removed from the configuration management system.

NEW QUESTION 8

You have been tasked with deploying a scalable distributed system using AWS OpsWorks. Your distributed system is required to scale on demand. As it is distributed, each node must hold a configuration file that includes the hostnames of the other instances within the layer. How should you configure AWS OpsWorks to manage scaling this application dynamically?

- A. Create a Chef Recipe to update this configuration file, configure your AWS OpsWorks stack to use custom cookbooks, and assign this recipe to the Configure Lifecycle Event of the specific layer.
- B. Update this configuration file by writing a script to poll the AWS OpsWorks service API for new instance
- C. Configure your base AMI to execute this script on Operating System startup.
- D. Create a Chef Recipe to update this configuration file, configure your AWS OpsWorks stack to use custom cookbooks, and assign this recipe to execute when instances are launched.
- E. Configure your AWS OpsWorks layer to use the AWS-provided recipe for distributed host configuration, and configure the instance hostname and file path parameters in your recipes settings.

Answer: A

Explanation:

Please check the following AWS DOCs which provides details on the scenario. Check the example of "configure".

? <https://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-events.html> You can use the Configure Lifecycle event

This event occurs on all of the stack's instances when one of the following occurs:

- An instance enters or leaves the online state.
- You associate an Elastic IP address with an instance or disassociate one from an instance.
- You attach an Elastic Load Balancing load balancer to a layer, or detach one from a layer. Ensure the Opswork layer uses a custom Cookbook

For more information on Opswork stacks, please refer to the below document link: from AWS

- http://docs.aws.amazon.com/opsworks/latest/userguide/welcome_classic.html

NEW QUESTION 9

You have a set of EC2 instances hosted in AWS. You have created a role named DemoRole and assigned that role to a policy, but you are unable to use that role with an instance. Why is this the case.

- A. You need to create an instance profile and associate it with that specific role.
- B. You are not able to associate an IAM role with an instance
- C. You won't be able to use that role with an instance unless you also create a user and associate it with that specific role
- D. You won't be able to use that role with an instance unless you also create a usergroup and associate it with that specific role.

Answer: A

Explanation:

An instance profile is a container for an IAM role that you can use to pass role information to an EC2 instance when the instance starts.

Option B is invalid because you can associate a role with an instance

Option C and D are invalid because using users or user groups is not a pre-requisite For more information on instance profiles, please visit the link:

- http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2-instance-profiles.html

NEW QUESTION 10

You are using Elastic Beanstalk to manage your application. You have a SQL script that needs to only be executed once per deployment no matter how many EC2 instances you have running. How can you do this?

- A. Use a "Container command" within an Elastic Beanstalk configuration file to execute the script, ensuring that the "leader only" flag is set to false.
- B. Use Elastic Beanstalk version and a configuration file to execute the script, ensuring that the "leader only" flag is set to true.
- C. Use a "Container command" within an Elastic Beanstalk configuration file to execute the script, ensuring that the "leader only" flag is set to true.
- D. Use a "leader command" within an Elastic Beanstalk configuration file to execute the script, ensuring that the "container only" flag is set to true.

Answer: C

Explanation:

You can use the container_commands key to execute commands that affect your application source code. Container commands run after the application and web

server have been set up and the application version archive has been extracted, but before the application version is deployed. Non- container commands and other customization operations are performed prior to the application source code being extracted.

You can use `leader_only` to only run the command on a single instance, or configure a test to only run the command when a test command evaluates to true. Leader-only container commands are only executed during environment creation and deployments, while other commands and server customization operations are performed every time an instance is provisioned or updated. Leader- only container commands are not executed due to launch configuration changes, such as a change in the AMI Id or instance type. For more information on customizing containers, please visit the below URL:
<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/customize-containers-ec2.html>

NEW QUESTION 10

You need to monitor specific metrics from your application and send real-time alerts to your Devops Engineer. Which of the below services will fulfil this requirement? Choose two answers

- A. Amazon CloudWatch
- B. Amazon Simple Notification Service
- C. Amazon Simple Queue Service
- D. Amazon Simple Email Service

Answer: AB

Explanation:

Amazon Cloud Watch monitors your Amazon Web Services (AWS) resources and the applications you run on AWS in real time. You can use Cloud Watch to collect and track metrics, which are variables you can measure for your resources and applications. Cloud Watch alarms send notifications or automatically make changes to the resources you are monitoring based on rules that you define.

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

• <http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/WhatIsCloudWatch.htm> | Amazon Cloud Watch uses Amazon SNS to send email. First, create and subscribe to an SNS topic.

When you create a CloudWatch alarm, you can add this SNS topic to send an email notification when the alarm changes state

For more information on AWS Cloudwatch and SNS, please refer to the below document link: from AWS

http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US_SetupSNS.html

NEW QUESTION 12

Your company has developed a web application and is hosting it in an Amazon S3 bucket configured for static website hosting. The application is using the AWS SDK for JavaScript in the browser to access data stored in an Amazon DynamoDB table. How can you ensure that API keys for access to your data in DynamoDB are kept secure?

- A. Create an Amazon S3 role in IAM with access to the specific DynamoDB tables, and assign it to the bucket hosting your website.
- B. Configure S3 bucket tags with your AWS access keys for your bucket hosing your website so that the application can query them for access.
- C. Configure a web identity federation role within IAM to enable access to the correct DynamoDB resources and retrieve temporary credentials.
- D. Store AWS keys in global variables within your application and configure the application to use these credentials when making requests.

Answer: C

Explanation:

With web identity federation, you don't need to create custom sign-in code or manage your own user identities. Instead, users of your app can sign in using a well-known identity provider (IdP) — such as Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible IdP, receive an authentication token, and then exchange that token for temporary security credentials in AWS that map to an IAM role with permissions to use the resources in your AWS account. Using an IdP helps you keep your AWS account secure, because you don't have to embed and distribute long- term security credentials with your application. For more information on Web Identity Federation, please refer to the below document link: from AWS

http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html

NEW QUESTION 14

You are using a configuration management system to manage your Amazon EC2 instances. On your Amazon EC2 Instances, you want to store credentials for connecting to an Amazon RDS MySQL DB instance. How should you securely store these credentials?

- A. Give the Amazon EC2 instances an IAM role that allows read access to a private Amazon S3 bucket
- B. Store a file with database credentials in the Amazon S3 bucket
- C. Have your configuration management system pull the file from the bucket when it is needed.
- D. Launch an Amazon EC2 instance and use the configuration management system to bootstrap the instance with the Amazon RDS DB credential
- E. Create an AMI from this instance.
- F. Store the Amazon RDS DB credentials in Amazon EC2 user data
- G. Import the credentials into the Instance on boot.
- H. Assign an IAM role to your Amazon EC2 instance, and use this IAM role to access the Amazon RDS DB from your Amazon EC2 instances.

Answer: D

Explanation:

Creating and Using an IAM Policy for IAM Database Access

To allow an IAM user or role to connect to your DB instance or DB cluster, you must create an IAM policy. After that you attach the policy to an IAM user or role. Note

To learn more about IAM policies, see Authentication and Access Control for Amazon RDS.

The following example policy allows an IAM user to connect to a DB instance using IAM database authentication.

Important

Don't confuse the `rds-db:` prefix with other Amazon RDS action prefixes that begin with `rds:`. You use the `rds-db:` prefix and the `rds-db:connect` action only for IAM database authentication. They aren't valid in any other context.

IAM Database Authentication for MySQL and Amazon Aurora

With Amazon RDS for MySQL or Aurora with MySQL compatibility, you can authenticate to your DB instance or DB cluster using AWS Identity and Access Management (IAM) database authentication. With this authentication method, you don't need to use a password when you connect to a DB instance. Instead, you use an authentication token.

An authentication token is a unique string of characters that Amazon RDS generates on request. Authentication tokens are generated using AWS Signature

Version 4. Each token has a lifetime of 15 minutes. You don't need to store user credentials in the database, because authentication is managed externally using IAM. You can also still use standard database authentication.

IAM database authentication provides the following benefits:

- Network traffic to and from the database is encrypted using Secure Sockets Layer (SSL).
- You can use IAM to centrally manage access to your database resources, instead of managing access individually on each DB instance or DB cluster.
- For applications running on Amazon EC2, you can use EC2 instance profile credentials to access the database instead of a password, for greater security.

For more information please refer to the below document link from AWS

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.IAMDBAuth.html>

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.IAMDBAuth.IAMPolicy.html>

You can use roles to delegate access to users, applications, or services that don't normally have access to your AWS resources. For example, you might want to grant users in your AWS account access to resources they don't usually have, or grant users in one AWS account access to resources in another account. Or you might want to allow a mobile app to use AWS resources, but not want to embed AWS keys within the app (where they can be difficult to rotate and where users can potentially extract them). Sometimes you want to give AWS access to users who already have identities defined outside of AWS, such as in your corporate directory. Or, you might want to grant access to your account to third parties so that they can perform an audit on your resources. For more information on IAM Roles, please refer to the below document link: from AWS

http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html

NEW QUESTION 16

You have an application hosted in AWS. This application was created using CloudFormation Templates and Autoscaling. Now your application has got a surge of users which is decreasing the performance of the application. As per your analysis, a change in the instance type to C3 would resolve the issue. Which of the below option can introduce this change while minimizing downtime for end users?

- A. Copy the old launch configuration, and create a new launch configuration with the C3 instance
- B. Update the Auto Scaling group with the new launch configuration
- C. Auto Scaling will then update the instance type of all running instances.
- D. Update the launch configuration in the AWS CloudFormation template with the new C3 instance type
- E. Add an UpdatePolicy attribute to the Auto Scaling group that specifies an AutoScalingRollingUpdate
- F. Run a stack update with the updated template.
- G. Update the existing launch configuration with the new C3 instance type
- H. Add an UpdatePolicy attribute to your Auto Scaling group that specifies an AutoScalingRollingUpdate in order to avoid downtime.
- I. Update the AWS CloudFormation template that contains the launch configuration with the new C3 instance type
- J. Run a stack update with the updated template, and Auto Scaling will then update the instances one at a time with the new instance type.

Answer: B

Explanation:

Ensure first that the cloudformation template is updated with the new instance type.

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.

Option A is invalid because this will cause an interruption to the users.

Option C is partially correct, but it does not have all the steps as mentioned in option B.

Option D is partially correct, but we need the AutoScalingRollingUpdate attribute to ensure a rolling update is performed.

For more information on AutoScaling Rolling updates please refer to the below link:

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

NEW QUESTION 19

What is web identity federation?

- A. Use of an identity provider like Google or Facebook to become an AWS IAM User.
- B. Use of an identity provider like Google or Facebook to exchange for temporary AWS security credentials.
- C. Use of AWS IAM User tokens to log in as a Google or Facebook user.
- D. Use STS service to create an user on AWS which will allow them to login from facebook or google app.

Answer: B

Explanation:

With web identity federation, you don't need to create custom sign-in code or manage your own user identities. Instead, users of your app can sign in using a well-known identity provider (IdP) — such as Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible IdP, receive an authentication token, and then exchange that token for temporary security credentials in AWS that map to an IAM role with permissions to use the resources in your AWS account. Using an IdP helps you keep your AWS account secure, because you don't have to embed and distribute long-term security credentials with your application. For more information on Web Identity federation please refer to the below link:

http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html

NEW QUESTION 21

You need to create a Route53 record automatically in CloudFormation when not running in production during all launches of a Template. How should you implement this?

- A. Use a Parameter for environment, and add a Condition on the Route53 Resource in the template to create the record only when environment is not production.
- B. Create two templates, one with the Route53 record value and one with a null value for the record
- C. Use the one without it when deploying to production.
- D. Use a Parameter for environment, and add a Condition on the Route53 Resource in the template to create the record with a null string when environment is production.
- E. Create two templates, one with the Route53 record and one without it
- F. Use the one without it when deploying to production.

Answer: A

Explanation:

The optional Conditions section includes statements that define when a resource is created or when a property is defined. For example, you can compare whether a value is equal to another value. Based on the result of that condition, you can conditionally create resources. If you have multiple conditions, separate them with commas.

You might use conditions when you want to reuse a template that can create resources in different contexts, such as a test environment versus a production environment. In your template, you can add an `EnvironmentType` input parameter, which accepts either `prod` or `test` as inputs. For the production environment, you might include Amazon EC2 instances with certain capabilities; however, for the test environment, you want to use reduced capabilities to save money. With conditions, you can define which resources are created and how they're configured for each environment type.

For more information on CloudFormation conditions please refer to the below link: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/conditions-section-structure.html>

NEW QUESTION 23

You have a development team that is continuously spending a lot of time rolling back updates for an application. They work on changes, and if the change fails, they spend more than 5-6h in rolling back the update. Which of the below options can help reduce the time for rolling back application versions.

- A. Use Elastic Beanstalk and re-deploy using Application Versions
- B. Use S3 to store each version and then re-deploy with Elastic Beanstalk
- C. Use CloudFormation and update the stack with the previous template
- D. Use OpsWorks and re-deploy using rollback feature.

Answer: A

Explanation:

Option B is invalid because Elastic Beanstalk already has the facility to manage various versions and you don't need to use S3 separately for this.

Option C is invalid because in CloudFormation you will need to maintain the versions. Elastic Beanstalk can do that automatically for you.

Option D is good for production scenarios and Elastic Beanstalk is great for development scenarios. AWS Beanstalk is the perfect solution for developers to maintain application versions.

With AWS Elastic Beanstalk, you can quickly deploy and manage applications in the AWS Cloud without worrying about the infrastructure that runs those applications. AWS Elastic Beanstalk reduces management complexity without restricting choice or control. You simply upload your application, and AWS Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring.

For more information on AWS Beanstalk please refer to the below link: <https://aws.amazon.com/documentation/elastic-beanstalk/>

NEW QUESTION 25

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 applications.

Answer: C

Explanation:

The first step in using Elastic Beanstalk is to create an application, which represents your web application in AWS. In Elastic Beanstalk an application serves as a container for the environments that run your web app, and versions of your web app's source code, saved configurations, logs and other artifacts that you create while using Elastic Beanstalk.

For more information on Applications, please refer to the below link: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/applications.html>

Deploying a new version of your application to an environment is typically a fairly quick process. The new source bundle is deployed to an instance and extracted, and the web container or application server picks up the new version and restarts if necessary. During deployment, your application might still become unavailable to users for a few seconds. You can prevent this by configuring your environment to use rolling deployments to deploy the new version to instances in batches. For more information on deployment, please refer to the below link: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-deploy-existing-version.html>

NEW QUESTION 26

You have decided to migrate your application to the cloud. You cannot afford any downtime. You want to gradually migrate so that you can test the application with a small percentage of users and increase over time. Which of these options should you implement?

- A. Use Direct Connect to route traffic to the on-premise location
- B. In DirectConnect, configure the amount of traffic to be routed to the on-premise location.
- C. Implement a Route 53 failover routing policy that sends traffic back to the on-premises application if the AWS application fails.
- D. Configure an Elastic Load Balancer to distribute the traffic between the on-premises application and the AWS application.
- E. Implement a Route 53 weighted routing policy that distributes the traffic between your on-premises application and the AWS application depending on weight.

Answer: D

Explanation:

Option A is incorrect because DirectConnect cannot control the flow of traffic.

Option B is incorrect because you want to split the percentage of traffic. Failover will direct all of the traffic to the backup servers.

Option C is incorrect because you cannot control the percentage distribution of traffic.

Weighted routing lets you associate multiple resources with a single domain name (example.com) or subdomain name (acme.example.com) and choose how much traffic is routed to each resource. This can be useful for a variety of purposes, including load balancing and testing new versions of software.

For more information on the Routing policy please refer to the below link: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

NEW QUESTION 28

For AWS Auto Scaling, what is the first transition state an instance enters after leaving steady state when scaling in due to health check failure or decreased load?

- A. Terminating
- B. Detaching
- C. Terminating:Wait

D. EnteringStandby

Answer: A

Explanation:

The below diagram shows the Lifecycle policy. When the scale-in happens, the first action is the Terminating action.

For more information on Autoscaling Lifecycle, please refer to the below link:
<http://docs.aws.amazon.com/autoscaling/latest/userguide/AutoScaingGroupLifecycle.html>

NEW QUESTION 31

You are using Chef in your data center. Which service is designed to let the customer leverage existing Chef recipes in AWS?

- A. AWS Elastic Beanstalk
- B. AWSOpsWorks
- C. AWS CloudFormation
- D. Amazon Simple Workflow Service

Answer: B

Explanation:

AWS OpsWorks is a configuration management service that uses Chef, an automation platform that treats server configurations as code. OpsWorks uses Chef to automate how servers are configured, deployed, and managed across your Amazon Elastic Compute Cloud (Amazon EC2) instances or on-premises compute environments. OpsWorks has two offerings, AWS Opsworks for Chef Automate, and AWS OpsWorks Stacks.

For more information on Opswork and SNS please refer to the below link:

- <https://aws.amazon.com/opsworks/>

NEW QUESTION 34

You need to deploy a Node.js application and do not have any experience with AWS. Which deployment method will be the simplest for you to deploy?

- A. AWS Elastic Beanstalk
- B. AWSCloudFormation
- C. AWS EC2
- D. AWSOpsWorks

Answer: A

Explanation:

With Elastic Beanstalk, you can quickly deploy and manage applications in the AWS Cloud without worrying about the infrastructure that runs those applications. AWS Elastic Beanstalk reduces management complexity without restricting choice or control. You simply upload your application, and Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring

For more information on Elastic beanstalk please refer to the below link:

- <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/Welcome.html>

NEW QUESTION 39

Which of these is not an intrinsic function in AWS CloudFormation?

- A. Fn::Equals
- B. Fn::If
- C. Fn::Not
- D. Fn::Parse

Answer: D

Explanation:

You can use intrinsic functions, such as Fn::If, Fn::Cquals, and Fn::Not, to conditionally create stack resources. These conditions are evaluated based on input parameters that you declare when you create or update a stack. After you define all your conditions, you can associate them with resources or resource properties in the Resources and Outputs sections of a template.

For more information on Cloud Formation template functions, please refer to the URL:

- <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference.html> and
- <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-conditions.html>

NEW QUESTION 41

There is a requirement to monitor API calls against your AWS account by different users and entities. There needs to be a history of those calls. The history of those calls are needed in bulk for later review. Which 2 services can be used in this scenario

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

Answer: C

Explanation:

You can use AWS CloudTrail to get a history of AWS API calls and related events for your account. This history includes calls made with the AWS Management Console, AWS Command Line Interface, AWS SDKs, and other AWS services. For more information on Cloudtrail, please visit the below URL:

- <http://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html>

Amazon Cloud Watch Cvents delivers a near real-time stream of system events that describe changes in Amazon Web Services (AWS) resources. Using simple rules that you can quickly set up, you can match events and route them to one or more target functions or streams. Cloud Watch Cvents becomes aware of operational changes as they occur. Cloud Watch Cvents responds to these operational changes and takes corrective action as necessary, by sending messages to respond to the environment, activating functions, making changes, and capturing state information. For more information on Cloud watch events, please visit the below URL:

- <http://docs.aws.amazon.com/AmazonCloudWatch/latest/events/WhatIsCloudWatchEvents.html>

NEW QUESTION 44

You have deployed a Cloudformation template which is used to spin up resources in your account. Which of the following status in Cloudformation represents a failure.

- A. UPDATE_COMPLETE_CLEANUP_IN_PROGRESS
- B. DELETE_COMPLETE
- C. ROLLBACK_IN_PROGRESS
- D. UPDATE_IN_PROGRESS

Answer: C

Explanation:

AWS Cloud Formation provisions and configures resources by making calls to the AWS services that are described in your template.

After all the resources have been created, AWS Cloud Formation reports that your stack has been created. You can then start using the resources in your stack. If stack creation fails, AWS CloudFormation rolls back your changes by deleting the resources that it created.

The below snapshot from Cloudformation shows what happens when there is an error in the stack creation.

For more information on how Cloud Formation works, please refer to the below link: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-what-is-howdoesitwork.html>

NEW QUESTION 49

You run accounting software in the AWS cloud. This software needs to be online continuously during the day every day of the week, and has a very static requirement for compute resources. You also have other, unrelated batch jobs that need to run once per day at anytime of your choosing. How should you minimize cost?

- A. Purchase a Heavy Utilization Reserved Instance to run the accounting software
- B. Turn it off after hour
- C. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.
- D. Purchase a Medium Utilization Reserved Instance to run the accounting software
- E. Turn it off after hour
- F. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.
- G. Purchase a Light Utilization Reserved Instance to run the accounting software
- H. Turn it off after hour
- I. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.
- J. Purchase a Full Utilization Reserved Instance to run the accounting software
- K. Turn it off after hour
- L. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.

Answer: A

Explanation:

Reserved Instances provide you with a significant discount compared to On-Demand Instance pricing.

Reserved Instances are not physical instances, but rather a

billing discount applied to the use of On-Demand Instances in your account. These On-Demand Instances must match certain attributes in order to benefit from the billing discount

For more information, please refer to the below link:

- <https://aws.amazon.com/about-aws/whats-new/2011/12/01/New-Amazon-CC2-Reserved-Instances-Options-Now-Available/>
- <https://aws.amazon.com/blogs/aws/reserved-instance-options-for-amazon-ec2/>
- <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-reserved-instances.html> Note:

It looks like these options are also no more available at present.

It looks like Convertible, Standard and scheduled are the new instance options. However the exams may still be referring to the old RIs.

<https://aws.amazon.com/ec2/pricing/reserved-instances/>

NEW QUESTION 53

You need the absolute highest possible network performance for a cluster computing application. You already selected homogeneous instance types supporting 10 gigabit enhanced networking, made sure that your workload was network bound, and put the instances in a placement group. What is the last optimization you can make?

- A. Use 9001 MTU instead of 1500 for Jumbo Frames, to raise packet body to packet overhead ratios.
- B. Segregate the instances into different peered VPCs while keeping them all in a placement group, so each one has its own Internet Gateway.
- C. Bake an AMI for the instances and relaunch, so the instances are fresh in the placement group and do not have noisy neighbors.
- D. Turn off SYN/ACK on your TCP stack or begin using UDP for higher throughput.

Answer: A

Explanation:

Jumbo frames allow more than 1500 bytes of data by increasing the payload size per packet, and thus increasing the percentage of the packet that is not packet overhead. Fewer packets are needed to send the same amount of usable data. However, outside of a given AWS region (CC2-Classic), a single VPC, or a VPC peering

connection, you will experience a maximum path of 1500 MTU. VPN connections and traffic sent over an Internet gateway are limited to 1500 MTU. If packets are over

1500 bytes, they are fragmented, or they are dropped if the Don't Fragment flag is set in the IP header.

For more information on Jumbo Frames, please visit the below URL:

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/network_mtu.htm#jumbo_frame_instances

NEW QUESTION 56

Your application's Auto Scaling Group scales up too quickly, too much, and stays scaled when traffic decreases. What should you do to fix this?

- A. Set a longer cooldown period on the Group, so the system stops overshooting the target capacity
- B. The issue is that the scaling system doesn't allow enough time for new instances to begin servicing requests before measuring aggregate load again.
- C. Calculate the bottleneck or constraint on the compute layer, then select that as the new metric, and set the metric thresholds to the bounding values that begin to affect response latency.
- D. Raise the CloudWatch Alarms threshold associated with your autoscaling group, so the scaling takes more of an increase in demand before beginning.
- E. Use larger instances instead of lots of smaller ones, so the Group stops scaling out so much and wasting resources as the OS level, since the OS uses a higher proportion of resources on smaller instances.

Answer: B

Explanation:

The ideal case is that the right metric is not being used for the scale up and down.

Option A is not valid because it mentions that the cooldown is not happening when the traffic decreases, that means the metric threshold for the scale down is not occurring in Cloudwatch

Option C is not valid because increasing the Cloudwatch alarm metric will not ensure that the instances scale down when the traffic decreases.

Option D is not valid because the question does not mention any constraints that points to the instance size. For an example on using custom metrics for scaling in and out, please follow the below link for a use case.

- <https://blog.powerupcloud.com/aws-autoscaling-based-on-database-query-custom-metrics-f396c16e5e6a>

NEW QUESTION 59

You are planning on using encrypted snapshots in the design of your AWS Infrastructure. Which of the following statements are true with regards to EBS Encryption

- A. Snapshotting an encrypted volume makes an encrypted snapshot; restoring an encrypted snapshot creates an encrypted volume when specified / requested.
- B. Snapshotting an encrypted volume makes an encrypted snapshot when specified / requested; restoring an encrypted snapshot creates an encrypted volume when specified / requested.
- C. Snapshotting an encrypted volume makes an encrypted snapshot; restoring an encrypted snapshot always creates an encrypted volume.
- D. Snapshotting an encrypted volume makes an encrypted snapshot when specified / requested; restoring an encrypted snapshot always creates an encrypted volume.

Answer: C

Explanation:

Amazon CBS encryption offers you a simple encryption solution for your CBS volumes without the need for you to build, maintain, and secure your own key management infrastructure. When you create an encrypted CBS volume and attach it to a supported instance type, the following types of data are encrypted:

- Data at rest inside the volume
- All data moving between the volume and the instance
- All snapshots created from the volume

Snapshots that are taken from encrypted volumes are automatically encrypted. Volumes that are created from encrypted snapshots are also automatically encrypted.

For more information on CBS encryption, please visit the below URL:

- <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/CBSEncryption.html>

NEW QUESTION 64

Your company wants to understand where cost is coming from in the company's production AWS account. There are a number of applications and services running at any given time. Without expending too much initial development time, how best can you give the business a good understanding of which applications cost the most per month to operate?

- A. Create an automation script which periodically creates AWS Support tickets requesting detailed intra-month information about your bill.
- B. Use custom CloudWatch Metrics in your system, and put a metric data point whenever cost is incurred.
- C. Use AWS Cost Allocation Tagging for all resources which support it.
- D. Use the Cost Explorer to analyze costs throughout the month.
- E. Use the AWS Price API and constantly running resource inventory scripts to calculate total price based on multiplication of consumed resources over time.

Answer: C

Explanation:

A tag is a label that you or AWS assigns to an AWS resource. Each tag consists of a key and a value. A key can have more than one value. You can use tags to organize your resources, and cost allocation tags to track your AWS costs on a detailed level. After you activate cost allocation tags, AWS uses the cost allocation tags to organize your resource costs on your cost allocation report, to make it easier

for you to categorize and track your AWS costs. AWS provides two types of cost allocation tags, an AWS-generated tag and user-defined tags. AWS defines, creates, and applies the AWS-generated tag for you, and you define, create, and apply user-defined tags. You must activate both types of tags separately before they can appear in Cost Explorer or on a cost allocation report.

For more information on Cost Allocation tags, please visit the below URL: <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-allocotags.html>

NEW QUESTION 68

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 overtime to the CTO.
- F. Use AWS Config with an SNS subscription on a Lambda, and insert these changes over time into a DynamoDB table
- G. Generate reports based on the contents of this table.

Answer: A

Explanation:

AWS CloudTrail is an AWS service that helps you enable governance, compliance, and operational and risk auditing of your AWS account. Actions taken by a user, role, or an AWS service are recorded as events in CloudTrail. Events include actions taken in the AWS Management Console, AWS Command Line Interface, and AWS SDKs and APIs.

Visibility into your AWS account activity is a key aspect of security and operational best practices. You can use CloudTrail to view, search, download, archive, analyze, and respond to account activity across your AWS infrastructure. You can identify who or what took which action, what resources were acted upon, when the event occurred, and other details to help you analyze and respond to activity in your AWS account.

For more information on Cloudtrail, please visit the below URL:

- <http://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html>

NEW QUESTION 71

You meet once per month with your operations team to review the past month's data. During the meeting, you realize that 3 weeks ago, your monitoring system which pings over HTTP from outside AWS recorded a large spike in latency on your 3-tier web service API. You use DynamoDB for the database layer, ELB, EBS, and EC2 for the business logic tier, and SQS, ELB, and EC2 for the presentation layer. Which of the following techniques will NOT help you figure out what happened?

- A. Check your CloudTrail log history around the spike's time for any API calls that caused slowness.
- B. Review CloudWatch Metrics for one minute interval graphs to determine which components slowed the system down.
- C. Review your ELB access logs in S3 to see if any ELBs in your system saw the latency.
- D. Analyze your logs to detect bursts in traffic at that time.

Answer: B

Explanation:

The Cloudwatch metric retention is as follows. If the data points are of a one minute interval, then the graphs will not be available in Cloudwatch

- Data points with a period of less than 60 seconds are available for 3 hours. These data points are high-resolution custom metrics.
- Data points with a period of 60 seconds (1 minute) are available for 15 days
- Data points with a period of 300 seconds (5 minute) are available for 63 days
- Data points with a period of 3600 seconds (1 hour) are available for 455 days (15 months) For more information on Cloudwatch metrics, please visit the below URL:
- http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/cloudwatch_concepts.html

NEW QUESTION 73

You need your CI to build AMIs with code pre-installed on the images on every new code push. You need to do this as cheaply as possible. How do you do this?

- A. Bid on spot instances just above the asking price as soon as new commits come in, perform all instance configuration and setup, then create an AMI based on the spot instance.
- B. Have the CI launch a new on-demand EC2 instance when new commits come in, perform all instance configuration and setup, then create an AMI based on the on-demand instance.
- C. Purchase a Light Utilization Reserved Instance to save money on the continuous integration machine
- D. Use these credits whenever you create AMIs on instances.
- E. When the CI instance receives commits, attach a new EBS volume to the CI machine
- F. Perform all setup on this EBS volume so you don't need

Answer: A

Explanation:

Amazon EC2 Spot instances allow you to bid on spare Amazon EC2 computing capacity. Since Spot instances are often available at a discount compared to On-Demand pricing, you can significantly reduce the cost of running your applications, grow your application's compute capacity and throughput for the same budget, and enable new types of cloud computing applications.

For more information on Spot Instances, please visit the below URL: <https://aws.amazon.com/ec2/spot/>

NEW QUESTION 78

You have an Autoscaling Group configured to launch EC2 Instances for your application. But you notice that the Autoscaling Group is not launching instances in the right proportion. In fact instances are being launched too fast. What can you do to mitigate this issue? Choose 2 answers from the options given below

- A. Adjust the cooldown period set for the Autoscaling Group
- B. Set a custom metric which monitors a key application functionality for the scale-in and scale-out process.
- C. Adjust the CPU threshold set for the Autoscaling scale-in and scale-out process.
- D. Adjust the Memory threshold set for the Autoscaling scale-in and scale-out process.

Answer: AB

Explanation:

The Auto Scaling cooldown period is a configurable setting for your Auto Scaling group that helps to ensure that Auto Scaling doesn't launch or terminate additional instances before the previous scaling activity takes effect.

For more information on the cool down period, please refer to the below link:

- <http://docs.aws.amazon.com/autoscaling/latest/userguide/Cooldown.html>

Also it is better to monitor the application based on a key feature and then trigger the scale-in and scale-out feature accordingly. In the question, there is no mention of CPU or memory causing the issue.

NEW QUESTION 80

There is a requirement for a vendor to have access to an S3 bucket in your account. The vendor already has an AWS account. How can you provide access to the vendor on this bucket.

- A. Create a new IAM user and grant the relevant access to the vendor on that bucket.
- B. Create a new IAM group and grant the relevant access to the vendor on that bucket.
- C. Create a cross-account role for the vendor account and grant that role access to the S3 bucket.
- D. Create an S3 bucket policy that allows the vendor to read from the bucket from their AWS account.

Answer: C

Explanation:

The AWS documentation mentions

You share resources in one account with users in a different account. By setting up cross-account access in this way, you don't need to create individual IAM users in each account. In addition, users don't have to sign out of one account and sign into another in order to access resources that are in different AWS accounts. After configuring the role, you see how to use the role from the AWS Management Console, the AWS CLI, and the API

For more information on Cross Account Roles Access, please refer to the below link:

- http://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 84

When building a multicontainer Docker platform using Elastic Beanstalk, which of the following is required

- A. DockerFile to create custom images during deployment
- B. Prebuilt Images stored in a public or private online image repository.
- C. Kubernetes to manage the docker containers.
- D. RedHat Openshift to manage the docker containers.

Answer: B

Explanation:

This is a special note given in the AWS Documentation for Multicontainer Docker platform for Elastic Beanstalk

Building custom images during deployment with a Dockerfile is not supported by the multicontainer Docker platform on Elastic Beanstalk. Build your images and deploy them to an online repository before creating an Elastic Beanstalk environment.

For more information on Multicontainer Docker platform for Elastic Beanstalk, please refer to the below link:

http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_ecs.html

NEW QUESTION 88

Which of the below services can be used to deploy application code content stored in Amazon S3 buckets, GitHub repositories, or Bitbucket repositories

- A. CodeCommit
- B. CodeDeploy
- C. S3Lifecycle
- D. Route53

Answer: B

Explanation:

The AWS documentation mentions

AWS CodeDeploy is a deployment service that automates application deployments to Amazon EC2 instances or on-premises instances in your own facility.

For more information on Code Deploy please refer to the below link:

- <http://docs.aws.amazon.com/codedeploy/latest/userguide/welcome.html>

NEW QUESTION 92

Which of the following is not a rolling type update which is present for Configuration Updates when it comes to the Elastic Beanstalk service

- A. Rolling based on Health
- B. Rolling based on Instances
- C. Immutable
- D. Rolling based on time

Answer: B

Explanation:

When you go to the configuration of your Elastic Beanstalk environment, below are the updates that are possible

The AWS Documentation mentions

- 1) With health-based rolling updates. Elastic Beanstalk waits until instances in a batch pass health checks before moving on to the next batch.
- 2) For time-based rolling updates, you can configure the amount of time that Elastic Beanstalk waits after completing the launch of a batch of instances before moving on to the next batch. This pause time allows your application to bootstrap and start serving requests.
- 3) Immutable environment updates are an alternative to rolling updates that ensure that configuration changes that require replacing instances are applied efficiently and safely. If an immutable environment update fails, the rollback process requires only terminating an Auto Scaling group. A failed rolling update, on the other hand, requires performing an additional rolling update to roll back the changes.

For more information on Rolling updates for Elastic beanstalk configuration updates, please visit the below URL:

- <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.ro11ingupdates.html>

NEW QUESTION 97

Your company has a set of resources hosted in AWS. Your IT Supervisor is concerned with the costs being incurred by the resources running in AWS and wants to optimize on the costs as much as possible. Which of the following ways could help achieve this efficiently? Choose 2 answers from the options given below.

- A. Create Cloudwatch alarms to monitor underutilized resources and either shutdown or terminate resources which are not required.
- B. Use the Trusted Advisor to see underutilized resources
- C. Create a script which monitors all the running resources and calculates the costs accordingly
- D. The analyze those resources accordingly and see which can be optimized.
- E. Create Cloudwatch logs to monitor underutilized resources and either shutdown or terminate resources which are not required.

Answer: AB

Explanation:

You can use Cloudwatch alarms to see if resources are below a threshold for long periods of time. If so you can take the decision to either stop them or to terminate the resources.

For more information on Cloudwatch alarms, please visit the below URL:

- [<http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/AlarmThatSendsEmail.html>](http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/AlarmThatSendsEmail.html)

In the Trusted Advisor, when you enable the Cost optimization section, you will get all sorts of checks which can be used to optimize the costs of your AWS resources.

For more information on the Trusted Advisor, please visit the below URL:

- <https://aws.amazon.com/premiumsupport/trustedadvisor/>

NEW QUESTION 100

Which of the following design strategies is ideal when designing loosely coupled systems. Choose 2 answers from the options given below

- A. Having the web and worker roles running on the same set of EC2 Instances
- B. Having the web and worker roles running on separate EC2 Instances
- C. Using SNS to establish communication between the web and worker roles

D. Using SQS to establish communication between the web and worker roles

Answer: BD

Explanation:

The below diagram shows the ideal design which uses SQS and separate environments for web and worker processes. The SQS queue manages the communication between the web and worker roles.

One example is the way Elastic beanstalk manages worker environments. For more information on this, please visit the below URL:
? <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.html>

NEW QUESTION 103

Your application is having a very high traffic, so you have enabled autoscaling in multi availability zone to suffice the needs of your application but you observe that one of the availability zone is not receiving any traffic. What can be wrong here?

- A. Autoscaling only works for single availability zone
- B. Autoscaling can be enabled for multi AZ only in north Virginia region
- C. Availability zone is not added to Elastic load balancer
- D. Instances need to manually added to availability zone

Answer: C

Explanation:

When you add an Availability Zone to your load balancer. Elastic Load Balancing creates a load balancer node in the Availability Zone. Load balancer nodes accept traffic from clients and forward requests to the healthy registered instances in one or more Availability Zones.

For more information on adding AZ's to CLB, please refer to the below URL:
<http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/enable-disable-az.html>

NEW QUESTION 107

Your company has an on-premise Active Directory setup in place. The company has extended their footprint on AWS, but still want to have the ability to use their on-premise Active Directory for authentication. Which of the following AWS services can be used to ensure that AWS resources such as AWS Workspaces can continue to use the existing credentials stored in the on-premise Active Directory.

- A. Use the Active Directory service on AWS
- B. Use the AWS Simple AD service
- C. Use the Active Directory connector service on AWS
- D. Use the ClassicLink feature on AWS

Answer: C

Explanation:

The AWS Documentation mentions the following

AD Connector is a directory gateway with which you can redirect directory requests to your on-premises Microsoft Active Directory without caching any information in the cloud. AD Connector comes in two sizes, small and large. A small AD Connector is designed for smaller organizations of up to 500 users. A large AD Connector can support larger organizations of up to 5,000 users.

For more information on the AD connector, please refer to the below URL: <http://docs.aws.amazon.com/directoryservice/latest/admin->

[guide/directory_ad_connector.html](#)

NEW QUESTION 108

You have a legacy application running that uses an m4.large instance size and cannot scale with Auto Scaling, but only has peak performance 5% of the time. This is a huge waste of resources and money so your Senior Technical Manager has set you the task of trying to reduce costs while still keeping the legacy application running as it should. Which of the following would best accomplish the task your manager has set you? Choose the correct answer from the options below

- A. Use a T2burstable performance instance.
- B. Use a C4.large instance with enhanced networking.
- C. Use two t2.nano instances that have single Root I/O Visualization.
- D. Use t2.nano instance and add spot instances when they are required.

Answer: A

Explanation:

The aws documentation clearly indicates using T2 CC2 instance types for those instances which don't use CPU that often.

T2

T2 instances are Burstable Performance Instances that provide a baseline level of CPU performance with the ability to burst above the baseline.

T2 Unlimited instances can sustain high CPU performance for as long as a workload needs it. For most general-purpose workloads, T2 Unlimited instances will provide ample performance without any additional charges. If the instance needs to run at higher CPU utilization for a prolonged period, it can also do so at a flat additional charge of 5 cents per vCPU-hour.

The baseline performance and ability to burst are governed by CPU Credits. T2 instances receive CPU Credits continuously at a set rate depending on the instance size, accumulating CPU Credits when they are idle, and consuming CPU credits when they are active. T2 instances are a good choice for a variety of general-purpose workloads including micro-services, low-latency interactive applications, small and medium databases, virtual desktops, development, build and stage environments, code repositories, and product prototypes. For more information see Burstable Performance Instances.

For more information on F_C2 instance types please see the below link: <https://aws.amazon.com/ec2/instance-types/>

NEW QUESTION 111

The company you work for has a huge amount of infrastructure built on AWS. However there has been some concerns recently about the security of this infrastructure, and an external auditor has been given the task of running a thorough check of all of your company's AWS assets. The auditor will be in the USA while your company's infrastructure resides in the Asia Pacific (Sydney) region on AWS. Initially, he needs to check all of your VPC assets, specifically, security groups and NACLs You have been assigned the task of providing the auditor with a login to be able to do this. Which of the following would be the best and most secure solution to provide the auditor with so he can begin his initial investigations? Choose the correct answer from the options below

- A. Create an 1AM user tied to an administrator rol
- B. Also provide an additional level of security with MFA.
- C. Give him root access to your AWS Infrastructure, because he is an auditor he will need access to every service.
- D. Create an 1AM user who will have read-only access to your AWS VPC infrastructure and provide the auditor with those credentials.
- E. Create an 1AM user with full VPC access but set a condition that will not allow him to modify anything if the request is from any IP other than his own.

Answer: C

Explanation:

Generally you should refrain from giving high level permissions and give only the required permissions. In this case option C fits well by just providing the relevant access which is required.

For more information on 1AM please see the below link:

- <https://aws.amazon.com/iam/>

NEW QUESTION 113

You're building a mobile application game. The application needs permissions for each user to communicate and store data in DynamoDB tables. What is the best method for granting each mobile device that installs your application to access DynamoDB tables for storage when required? Choose the correct answer from the options below

- A. During the install and game configuration process, have each user create an 1AM credential and assign the 1AM user to a group with proper permissions to communicate with DynamoDB.
- B. Create an 1AM group that only gives access to your application and to the DynamoDB table
- C. Then, when writing to DynamoDB, simply include the unique device ID to associate the data with that specific user.
- D. Create an 1AM role with the proper permission policy to communicate with the DynamoDB tabl
- E. Use web identity federation, which assumes the 1AM role using AssumeRoleWithWebIdentity, when the user signs in, granting temporary security credentials using STS.
- F. Create an Active Directory server and an AD user for each mobile application use
- G. When the user signs in to the AD sign-on, allow the AD server to federate using SAML 2.0 to 1AM and assign a role to the AD user which is the assumed with AssumeRoleWithSAML

Answer: C

Explanation:

Answer - C

For access to any AWS service, the ideal approach for any application is to use Roles. This is the first preference.

For more information on 1AM policies please refer to the below link:

http://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html

Next for any web application, you need to use web identity federation. Hence option D is the right option. This along with the usage of roles is highly stressed in the aws documentation.

The AWS documentation mentions the following

When developing a web application it is recommend not to embed or distribute long-term AWS credentials with apps that a user downloads to a device, even in an encrypted store. Instead, build your app so that it requests temporary AWS security credentials dynamically when needed using web identity federation. The supplied temporary credentials map to an AWS role that has only the permissions needed to perform the tasks required by the mobile app.

For more information on web identity federation please refer to the below link: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html

NEW QUESTION 114

You are in charge of creating a CloudFormation template that will be used to spin our resources on demand for your Devops team. The requirement is that this cloudformation template should be able to spin up resources in different regions. Which of the following aspects of Cloudformation templates can help you design the template to spin up resources based on the region.

- A. Use mappings section in the Cloudformation template, so that based on the relevant region, the relevant resource can be spinned up.
- B. Use the outputs section in the Cloudformation template, so that based on the relevant region, the relevant resource can be spinned up.
- C. Use the parameters section in the Cloudformation template, so that based on the relevant region, the relevant resource can be spinned up.
- D. Use the metadata section in the Cloudformation template, so that based on the relevant region, the relevant resource can be spinned up.

Answer: A

Explanation:

The AWS Documentation mentions

The optional Mappings section matches a key to a corresponding set of named values. For example, if you want to set values based on a region, you can create a mapping that uses the region name as a key and contains the values you want to specify for each specific region. You use the Fn::FindInMap intrinsic function to retrieve values in a map.

For more information on mappings please refer to the below link:

? <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/mappings-section-structure.html>

NEW QUESTION 116

A company has developed a Ruby on Rails content management platform. Currently, OpsWorks with several stacks for dev, staging, and production is being used to deploy and manage the application. Now the company wants to start using Python instead of Ruby. How should the company manage the new deployment? Choose the correct answer from the options below

- A. Update the existing stack with Python application code and deploy the application using the deploy life-cycle action to implement the application code.
- B. Create a new stack that contains a new layer with the Python code
- C. To cut over to the new stack the company should consider using Blue/Green deployment
- D. Create a new stack that contains the Python application code and manage separate deployments of the application via the secondary stack using the deploy lifecycle action to implement the application code.
- E. Create a new stack that contains the Python application code and manages separate deployments of the application via the secondary stack.

Answer: B

Explanation:

Blue/green deployment is a technique for releasing applications by shifting traffic between two identical environments running different versions of the application. Blue/green deployments can mitigate common risks associated with deploying software, such as downtime and rollback capability

Please find the below link on a white paper for blue green deployments

• https://d03wsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 118

Which of the following features of the Autoscaling Group ensures that additional instances are neither launched or terminated before the previous scaling activity takes effect

- A. Termination policy
- B. Cool down period
- C. Ramp up period
- D. Creation policy

Answer: B

Explanation:

The AWS documentation mentions

The Auto Scaling cooldown period is a configurable setting for your Auto Scaling group that helps to ensure that Auto Scaling doesn't launch or terminate additional

instances before the previous scaling activity takes effect. After the Auto Scaling group dynamically scales using a simple scaling policy. Auto Scaling waits for the cooldown period to complete before resuming scaling activities. When you manually scale your Auto Scaling group, the default is not to wait for the cooldown period,

but you can override the default and honor the cooldown period. If an instance becomes unhealthy.

Auto Scaling does not wait for the cooldown period to complete before replacing the unhealthy instance

For more information on the Cool down period, please refer to the below URL:

• <http://docs.ws.amazon.com/autoscaling/latest/userguide/Cooldown.html>

NEW QUESTION 119

There is a requirement for an application hosted on a VPC to access the On-premise LDAP server. The VPC and the On-premise location are connected via an IPsec VPN. Which of the below are the right options for the application to authenticate each user. Choose 2 answers from the options below

- A. Develop an identity broker that authenticates against 1AM security Token service to assume a 1AM role in order to get temporary AWS security credentials The application calls the identity broker to get AWS temporary security credentials.
- B. The application authenticates against LDAP and retrieves the name of an 1AM role associated with the user
- C. The application then calls the 1AM Security Token Service to assume that 1AM role
- D. The application can use the temporary credentials to access any AWS resources.
- E. Develop an identity broker that authenticates against LDAP and then calls 1AM Security Token Service to get 1AM federated user credential
- F. The application calls the identity broker to get 1AM federated user credentials with access to the appropriate AWS service.
- G. The application authenticates against LDAP the application then calls the AWS identity and Access Management (IAM) Security service to log in to 1AM using the LDAP credentials the application can use the 1AM temporary credentials to access the appropriate AWS service.

Answer: BC

Explanation:

When you have the need for an in-premise environment to work with a cloud environment, you would normally have 2 artefacts for authentication purposes

- An identity store - So this is the on-premise store such as Active Directory which stores all the information for the user's and the groups they belong to.
- An identity broker - This is used as an intermediate agent between the on-premise location and the cloud environment. In Windows you have a system known as Active Directory Federation services to provide this facility.

Hence in the above case, you need to have an identity broker which can work with the identity store and the Security Token service in AWS. An example diagram of how this works from the AWS documentation is given below.

For more information on federated access, please visit the below link: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_common-scenarios_federated-users.html

NEW QUESTION 124

You are planning on using AWS Code Deploy in your AWS environment. Which of the below features of AWS Code Deploy can be used to Specify scripts to be run on each instance at various stages of the deployment process

- A. AppSpecfile
- B. CodeDeployfile
- C. Configfile
- D. Deploy file

Answer: A

Explanation:

The AWS Documentation mentions the following on AWS Code Deploy

An application specification file (AppSpec file), which is unique to AWS CodeDeploy, is a YAML- formatted file used to:

Map the source files in your application revision to their destinations on the instance. Specify custom permissions for deployed files.

Specify scripts to be run on each instance at various stages of the deployment process. For more information on AWS CodeDeploy, please refer to the URL:

<http://docs.aws.amazon.com/codedeploy/latest/userguide/application-specification-files.html>

NEW QUESTION 126

Of the 6 available sections on a Cloud Formation template (Template Description Declaration, Template Format Version Declaration, Parameters, Resources, Mappings, Outputs), which is the only one required for a CloudFormation template to be accepted? Choose an answer from the options below

- A. Parameters
- B. Template Declaration
- C. Mappings
- D. Resources

Answer: D

Explanation:

If you refer to the documentation, you will see that Resources is the only mandatory field

Specifies the stack resources and their properties, such as an Amazon Elastic Compute Cloud instance or an Amazon Simple Storage Service bucket.

For more information on CloudFormation templates, please refer to the below link:

- <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-anatomy.html>

NEW QUESTION 127

An organization is planning to use AWS for their Production Rollout. The organizations wants to implement automation for deployment, such that it will automatically create a LAMP stack, deploy an RDS MySQLDB instance, download the latest PHP installable from S3 and set up the ELB. Which of the below mentioned AWS services meets the requirement for making an orderly deployment of the software?

- A. AWS Elastic Beanstalk
- B. AWS Cloudfront
- C. AWS Cloudformation
- D. AWS DevOps

Answer: C

Explanation:

When you want to automate deployment, the automatic choice is Cloudformation. Below is the excerpt from AWS on cloudformation.

AWS Cloud Formation gives developers and systems administrators an easy way to create and manage a collection of related AWS resources, provisioning updating them in an orderly and predictable fashion.

You can use AWS Cloud Formation's sample templates or create your own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run your application. You don't need to figure out the order for provisioning AWS services or the subtleties of making those dependencies work. Cloud Formation takes care of this for you. After the AWS resources are deployed, you can modify and update them in a controlled and predictable way, in effect applying version control to your AWS infrastructure the same way you do with your software For more information on Cloud Formation, please visit the link:

- <https://aws.amazon.com/cloudformation/> As per AWS,

"AWS Elastic Beanstalk provides support for running Amazon Relational Database Service (Amazon RDS) instances in your Elastic Beanstalk environment. This works great for development and testing environments. However, it isn't ideal for a production environment because it ties the lifecycle of the database instance to the lifecycle of your application's environment."

- <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/AWSHowTo.RDS.html>

NEW QUESTION 128

You are currently using SNS to pass messages to EC2 Instances. You need to pass messages which are greater than 5 MB in size. Which of the following can help you accomplish this.

- A. Use Kinesis as a buffer stream for message bodies
- B. Store the checkpoint id for the placement in the Kinesis Stream in SQS.
- C. Use the Amazon SQS Extended Client Library for Java and Amazon S3 as a storage mechanism for message bodies
- D. */
- E. Use SQS's support for message partitioning and multi-part uploads on Amazon S3.
- F. Use AWS EFS as a shared pool storage medium
- G. Store filesystem pointers to the files on disk in the SQS message bodies.

Answer: B

Explanation:

The AWS documentation mentions the following

You can manage Amazon SQS messages with Amazon S3. This is especially useful for storing and consuming messages with a message size of up to 2 GB. To manage

Amazon SQS messages with Amazon S3, use the Amazon SQS Extended Client Library for Java. Specifically, you use this library to:

Specify whether messages are always stored in Amazon S3 or only when a message's size exceeds 256 KB.

Send a message that references a single message object stored in an Amazon S3 bucket. Get the corresponding message object from an Amazon S3 bucket.

Delete the corresponding message object from an Amazon S3 bucket. For more information on SQS and sending larger messages please visit the link

NEW QUESTION 133

Which of the following run command types are available for OpsWorks stacks? Choose 3 answers from the options given below.

- A. UpdateCustom Cookbooks
- B. Execute Recipes
- C. Configure
- D. Undeploy

Answer: ABC

NEW QUESTION 136

Your company owns multiple AWS accounts. There is currently one development and one production account. You need to grant access to the development team to an S3 bucket in the production account. How can you achieve this?

- A. Create an IAM user in the Production account that allows users from the Development account (the trusted account) to access the S3 bucket in the Production account.
- B. When creating the role, define the Development account as a trusted entity and specify a permissions policy that allows trusted users to update the S3 bucket.
- C. Use web identity federation with a third-party identity provider with AWS STS to grant temporary credentials and membership into the production IAM user.
- D. Create an IAM cross account role in the Production account that allows users from the Development account to access the S3 bucket in the Production account.

Answer: D

Explanation:

The AWS Documentation mentions the following on cross account roles

You can use AWS Identity and Access Management (IAM) roles and AWS Security Token Service (STS) to set up cross-account access between AWS accounts.

When you assume an IAM role in another AWS account to obtain cross-account access to services and resources in that account, AWS CloudTrail logs the cross-account activity. For more information on Cross account roles, please visit the below URL

- http://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 139

Explain what the following resource in a CloudFormation template does? Choose the best possible answer.

- A. Creates an SNS topic which allows SQS subscription endpoints to be added as a parameter on the template
- B. Creates an SNS topic that allows SQS subscription endpoints
- C. Creates an SNS topic and then invokes the call to create an SQS queue with a logical resource name of SQSQueue
- D. Creates an SNS topic and adds a subscription ARN endpoint for the SQS resource created under the logical name SQSQueue

Answer: D

Explanation:

The intrinsic function `Fn::GetAtt` returns the value of an attribute from a resource in the template. This has nothing to do with adding parameters (Option A is wrong) or allowing endpoints (Option B is wrong) or invoking relevant calls (Option C is wrong)

For more information on `Fn::GetAtt` function please refer to the below link

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-getatt.html>

NEW QUESTION 142

You are in charge of designing a CloudFormation template which deploys a LAMP stack. After deploying a stack, you see that the status of the stack is showing as `CREATE_COMPLETE`, but the Apache server is still not up and running and is experiencing issues while starting up. You want to ensure that the stack creation only shows the status of `CREATE_COMPLETE` after all resources defined in the stack are up and running. How can you achieve this?

Choose 2 answers from the options given below.

- A. Define a stack policy which defines that all underlying resources should be up and running before showing a status of `CREATE_COMPLETE`.
- B. Use lifecycle hooks to mark the completion of the creation and configuration of the underlying resource.
- C. Use the `CreationPolicy` to ensure it is associated with the EC2 Instance resource.
- D. Use the CFN helper scripts to signal once the resource configuration is complete.

Answer: CD

Explanation:

The AWS Documentation mentions

When you provision an Amazon EC2 instance in an AWS CloudFormation stack, you might specify additional actions to configure the instance, such as install software packages or bootstrap applications. Normally, CloudFormation proceeds with stack creation after the instance has been successfully created. However, you can use a `CreationPolicy` so that CloudFormation proceeds with stack creation only after your configuration actions are done. That way you'll know your applications are ready to go after stack creation succeeds.

For more information on the `CreationPolicy`, please visit the below URL <https://aws.amazon.com/blogs/devops/use-a-creationpolicy-to-wait-for-on-instance-configurations/>

NEW QUESTION 147

You are a DevOps Engineer in your company. You have been instructed to ensure there is an automated backup solution in place for EBS Volumes. These snapshots need to be retained only for a period of 20 days. How can you achieve this requirement in an efficient manner?

- A. Use the aws ec2 create-volume API to create a snapshot of the EBS Volume
- B. Use the describe-volumes API to see those snapshots which are greater than 20 days and then delete them accordingly using the delete-volume API call.
- C. Use Lifecycle policies to push the EBS Volumes to Amazon Glacie
- D. Then use further lifecycle policies to delete the snapshots after 20 days.
- E. Use Lifecycle policies to push the EBS Volumes to Amazon S3. Then use further lifecycle policies to delete the snapshots after 20 days.
- F. Use Amazon Data Lifecycle Manager to automate the process.

Answer: D

Explanation:

Use Amazon Data Lifecycle Manager (Amazon DLM) to automate the creation, retention, and deletion of snapshots taken to back up your Amazon EBS volumes. Automating snapshot management helps you to:

- Protect valuable data by enforcing a regular backup schedule. Retain backups as required by auditors or internal compliance.
- Reduce storage costs by deleting outdated backups.

For more information, please check the below AWS Docs:

- <https://docs.aws.amazon.com/AWSCC2/latest/UserGuide/snapshot-lifecycle.html>

NEW QUESTION 149

You are DevOps Engineer for a large organization. The company wants to start using CloudFormation templates to start building their resources in AWS. You are getting requirements for the templates from various departments, such as the networking, security, application etc. What is the best way to architect these CloudFormation templates.

- A. Use a single CloudFormation template, since this would reduce the maintenance overhead on the templates itself.
- B. Create separate logical templates, for example, a separate template for networking, security, application etc.
- C. Then nest the relevant templates.
- D. Consider using Elastic Beanstalk to create your environments since CloudFormation is not built for such customization.
- E. Consider using Opsworks to create your environments since CloudFormation is not built for such customization.

Answer: B

Explanation:

The AWS documentation mentions the following

As your infrastructure grows, common patterns can emerge in which you declare the same components in each of your templates. You can separate out these common components and create dedicated templates for them. That way, you can mix and match different templates but use nested stacks to create a single, unified stack. Nested stacks are stacks that create other stacks. To create nested stacks, use the AWS::

CloudFormation::Stack resource in your template to reference other templates.

For more information on CloudFormation best practices, please visit the below URL: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html>

NEW QUESTION 150

Which of the following tools for EC2 can be used to administer instances without the need to SSH or RDP into the instance.

- A. AWSConfig
- B. AWSCodePipeline
- C. RunCommand
- D. EC2Config

Answer: C

Explanation:

You can use Run Command from the Amazon EC2 console to configure instances without having to login to each instance

For more information on the Run Command, please visit the below URL:

- <http://docs.aws.amazon.com/systems-manager/latest/userguide/rc-console.html>

NEW QUESTION 152

You are trying to debug the creation of CloudFormation stack resources. Which of the following can be used to help in the debugging process?

Choose 2 answers from the options below

- A. Use CloudTrail to debug all the API calls sent by the CloudFormation stack.
- B. Use the AWS CloudFormation console to view the status of your stack.
- C. See the logs in the /var/log directory for Linux instances
- D. Use AWSConfig to debug all the API calls sent by the CloudFormation stack.

Answer: BC

Explanation:

The AWS Documentation mentions

Use the AWS CloudFormation console to view the status of your stack. In the console, you can view a list of stack events while your stack is being created, updated, or

deleted. From this list, find the failure event and then view the status reason for that event.

For Amazon EC2 issues, view the cloud-init and cfn logs. These logs are published on the Amazon EC2 instance in the /var/log/ directory. These logs capture processes and command outputs while AWS CloudFormation is setting up your instance. For Windows, view the EC2Configure service and cfn logs in %ProgramFiles%\Amazon\EC2ConfigureService and C:\cfn\log.

For more information on CloudFormation Troubleshooting, please visit the below URL:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/troubleshooting.html>

NEW QUESTION 157

Your team is responsible for an AWS Elastic Beanstalk application. The business requires that you move to a continuous deployment model, releasing updates to the application multiple times per day with zero downtime. What should you do to enable this and still be able to roll back almost immediately in an emergency to

the previous version?

- A. Enable rolling updates in the Elastic Beanstalk environment, setting an appropriate pause time for application startup.
- B. Create a second Elastic Beanstalk environment running the new application version, and swap the environment CNAMEs.
- C. Develop the application to poll for a new application version in your code repository; download and install to each running Elastic Beanstalk instance.
- D. Create a second Elastic Beanstalk environment with the new application version, and configure the old environment to redirect clients, using the HTTP 301 response code, to the new environment

Answer: B

Explanation:

The AWS Documentation mentions the below

Because Elastic Beanstalk performs an in-place update when you update your application versions, your application may become unavailable to users for a short period of time. It is possible to avoid this downtime by performing a blue/green deployment, where you deploy the new version to a separate environment, and then swap CNAMEs of the two environments to redirect traffic to the new version instantly. For more information on Elastic Beanstalk swap URL, please see the below link:

- <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAME-Change.html>

NEW QUESTION 160

Which of the following can be configured as targets for Cloudwatch Events. Choose 3 answers from the options given below

- A. Amazon EC2 Instances
- B. AWS Lambda Functions
- C. Amazon CodeCommit
- D. Amazon ECS Tasks

Answer: ABD

Explanation:

The AWS Documentation mentions the below

You can configure the following AWS services as targets for CloudWatch Events

For more information on Cloudwatch events, please see the below link:

- <http://docs.aws.amazon.com/AmazonCloudWatch/latest/events/WhatIsCloudWatchEvents.html>

NEW QUESTION 165

When you implement a lifecycle hook in Autoscaling, by default what is the time limit in which the instance will be in a pending state.

- A. 60 seconds
- B. 5 minutes
- C. 60 minutes
- D. 120 minutes

Answer: C

Explanation:

The AWS Documentation mentions

By default, the instance remains in a wait state for one hour, and then Auto Scaling continues the launch or terminate process (Pending: Proceed or Terminating: Proceed). If you need more time, you can restart the timeout period by recording a heartbeat. If you finish before the timeout period ends, you can complete the lifecycle action, which continues the launch or termination process.

For more information on Autoscaling lifecycle hooks, please see the below link:

- <http://docs.aws.amazon.com/autoscaling/latest/userguide/lifecycle-hooks.html>

NEW QUESTION 170

You are a DevOps Engineer for your company. You are planning on using Cloudwatch for monitoring the resources hosted in AWS. Which of the following can you do with Cloudwatch logs ideally. Choose 3 answers from the options given below

- A. Stream the log data to Amazon Kinesis for further processing
- B. Send the log data to AWS Lambda for custom processing
- C. Stream the log data into Amazon Elasticsearch for any search analysis required.
- D. Send the data to SQS for further processing.

Answer: ABC

Explanation:

Amazon Kinesis can be used for rapid and continuous data intake and aggregation. The type of data used includes IT infrastructure log data, application logs, social media, market data feeds, and web clickstream data. Amazon Lambda is a web service which can be used to do serverless computing of the logs which are published by Cloudwatch logs. Amazon Elasticsearch Service makes it easy to deploy, operate, and scale Elasticsearch for log analytics, full text search, application monitoring, and more.

For more information on Cloudwatch logs, please see the below link: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html>

NEW QUESTION 174

Which of the following tools is available to send log data from EC2 Instances.

- A. CloudWatch LogsAgent
- B. CloudWatchAgent
- C. Logsconsole.
- D. LogsStream

Answer: A

Explanation:

The AWS Documentation mentions the following

The CloudWatch Logs agent provides an automated way to send log data to Cloud Watch Logs from Amazon EC2 instances. The agent is comprised of the following components:

A plug-in to the AWS CLI that pushes log data to CloudWatch Logs.

A script (daemon) that initiates the process to push data to CloudWatch Logs.

Acron job that ensures that the daemon is always running. For more information on Cloudwatch logs Agent, please see the below link:

<http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/AgentReference.html>

NEW QUESTION 178

You work as a Devops Engineer for your company. There are currently a number of environments hosted via Elastic beanstalk. There is a requirement to ensure to ensure that the rollback time for a new version application deployment is kept to a minimal. Which elastic beanstalk deployment method would fulfil this requirement ?

- A. Rollingwith additional batch
- B. AllatOnce
- C. Blue/Green
- D. Rolling

Answer: C

Explanation:

The below table from the AWS documentation shows that the least amount of time is spent in rollbacks when it comes to Blue Green deployments. This is because the only thing that needs to be done is for URL's to be swapped.

For more information on Elastic beanstalk deployment strategies, please visit the below URL: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.deploy-existing-version.html>

NEW QUESTION 179

By default in Opswork, how many application versions can you rollback up to?

- A. 1
- B. 2
- C. 3
- D. 4

Answer: D

Explanation:

The AWS Documentation mentions the following Restores the previously deployed app version. For example, if you have deployed the app three times and then run Rollback, the server will serve the app from the second deployment. If you run Rollback again, the server will serve the app from the first deployment. By default, AWS OpsWorks Stacks stores the five most recent deployments, which allows you to roll back up to four versions. If you exceed the number of stored versions, the command fails and leaves the oldest version in place.

For more information on Opswork app deployment, please visit the below URL: <http://docs.aws.amazon.com/opsworks/latest/userguide/workingapps-deploying.html>

NEW QUESTION 181

A company wants to create standard templates for deployment of their Infrastructure. Which AWS service can be used in this regard? Please choose one option.

- A. AmazonSimple Workflow Service
- B. AWSElastic Beanstalk
- C. AWSCloudFormation
- D. AWSOpsWorks

Answer: C

Explanation:

AWS Cloud Formation gives developers and systems administrators an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable fashion.

You can use AWS Cloud Formation's sample templates or create your own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run your application. You don't need to figure out the order for provisioning AWS services or the subtleties of making those dependencies work. Cloud Formation takes care of this for you. After the AWS resources are deployed, you can modify and update them in a controlled and predictable way, in effect applying version control to your AWS infrastructure the same way you do with your software. You can also visualize your templates as diagrams and edit them using a drag-and-drop interface with the AWS CloudFormation Designer.

For more information on Cloudformation, please visit the link:

- <https://aws.amazon.com/cloudformation/>

NEW QUESTION 186

The development team has developed a new feature that uses an AWS service and wants to test it from inside a staging VPC. How should you test this feature with the fastest turnaround time?

- A. Launch an Amazon Elastic Compute Cloud (EC2) instance in the staging VPC in response to a development request, and use configuration management to set up the application.
- B. Run any testing harnesses to verify application functionality and then use Amazon Simple Notification Service (SNS) to notify the development team of the results.
- C. Use an Amazon EC2 instance that frequently polls the version control system to detect the new feature, use AWS CloudFormation and Amazon EC2 user data to run any testing harnesses to verify application functionality and then use Amazon SNS to notify the development team of the results.
- D. Use an Elastic Beanstalk application that polls the version control system to detect the new feature, use AWS CloudFormation and Amazon EC2 user data to run any testing harnesses to verify application functionality and then use Amazon Kinesis to notify the development team of the results.
- E. Use AWS CloudFormation to launch an Amazon EC2 instance, use Amazon EC2 user data to run any testing harnesses to verify application functionality and then use Amazon Kinesis to notify the development team of the results.

Answer: A

Explanation:

Using Amazon Kinesis would just take more time in setup and would not be ideal to notify the relevant team in the shortest time possible.

Since the test needs to be conducted in the staging VPC, it is best to launch the EC2 in the staging VPC.

For more information on the Simple Notification service, please visit the link:

- <https://aws.amazon.com/sns/>

NEW QUESTION 190

Your application uses Amazon SQS and Auto Scaling to process background jobs. The Auto Scaling policy is based on the number of messages in the queue, with a maximum instance count of 100. Since the application was launched, the group has never scaled above 50. The Auto scaling group has now scaled to 100, the queue size is increasing and very few jobs are being completed. The number of messages being sent to the queue is at normal levels. What should you do to identify why the queue size is unusually high and to reduce it?

- A. Temporarily increase the AutoScaling group's desired value to 200. When the queue size has been reduced, reduce it to 50.
- B. Analyze the application logs to identify possible reasons for message processing failure and resolve the cause for failure.
- C. V
- D. Create additional Auto Scaling groups enabling the processing of the queue to be performed in parallel.
- E. Analyze CloudTrail logs for Amazon SQS to ensure that the instances Amazon EC2 role has permission to receive messages from the queue.

Answer: B

Explanation:

Here the best option is to look at the application logs and resolve the failure. You could be having a functionality issue in the application that is causing the messages to queue up and increase the fleet of instances in the AutoScaling group.

For more information on centralized logging system implementation in AWS, please visit this link: <https://aws.amazon.com/answers/logging/centralized-logging/>

NEW QUESTION 195

As part of your deployment pipeline, you want to enable automated testing of your AWS CloudFormation template. What testing should be performed to enable faster feedback while minimizing costs and risk? Select three answers from the options given below

- A. Use the AWS CloudFormation Validate Template to validate the syntax of the template.
- B. Use the AWS CloudFormation Validate Template to validate the properties of resources defined in the template.
- C. Validate the template's syntax using a general JSON parser.
- D. Validate the AWS CloudFormation template against the official XSD scheme definition published by Amazon Web Services.
- E. Update the stack with the template.
- F. If the template fails, rollback will return the stack and its resources to exactly the same state.
- G. When creating the stack, specify an Amazon SNS topic to which your testing system is subscribed.
- H. Your testing system runs tests when it receives notification that the stack is created or updated.

Answer: AEF

Explanation:

The AWS documentation mentions the following

The `aws cloudformation validate-template` command is designed to check only the syntax of your template. It does not ensure that the property values that you have specified for a resource are valid for that resource. Nor does it determine the number of resources that will exist when the stack is created.

To check the operational validity, you need to attempt to create the stack. There is no sandbox or test area for AWS CloudFormation stacks, so you are charged for the resources you create during testing. Option F is needed for notification.

For more information on CloudFormation template validation, please visit the link:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-validation-template.html>

NEW QUESTION 199

You set up a web application development environment by using a third party configuration management tool to create a Docker container that is run on local developer machines.

What should you do to ensure that the web application and supporting network storage and security infrastructure does not impact your application after you deploy into AWS for staging and production environments?

- A. Write a script using the AWS SDK or CLI to deploy the application code from version control to the local development environments, staging and production using AWS OpsWorks.
- B. Define an AWS CloudFormation template to place your infrastructure into version control and use the same template to deploy the Docker container into Elastic Beanstalk for staging and production.

- C. Because the application is inside a Docker container, there are no infrastructure differences to be taken into account when moving from the local development environments to AWS for staging and production.
- D. Define an AWS CloudFormation template for each stage of the application deployment lifecycle - development, staging and production - and have tagging in each template to define the environment.

Answer: B

Explanation:

Elastic Beanstalk supports the deployment of web applications from Docker containers. With Docker containers, you can define your own runtime environment. You can choose your own platform, programming language, and any application dependencies (such as package managers or tools), that aren't supported by other platforms. Docker containers are self-contained and include all the configuration information and software your web application requires to run. By using Docker with Elastic Beanstalk, you have an infrastructure that automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring.

This seems to be more appropriate than Option D.

? https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker.html

For more information on CloudFormation best practices, please visit the link: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html>

NEW QUESTION 202

Which of the following services allows you to easily run and manage Docker-enabled applications across a cluster of Amazon EC2 instances

- A. Elastic bean stalk
- B. ElasticContainer service
- C. Opswork
- D. Cloudwatch

Answer: B

Explanation:

The AWS documentation provides the following information

Amazon EC2 Container Service (ECS) allows you to easily run and manage Docker-enabled applications across a cluster of Amazon EC2 instances. Applications packaged as containers locally will deploy and run in the same way as containers managed by Amazon ECS. Amazon ECS eliminates the need to install, operate, and scale your own cluster management infrastructure, and allows you to schedule Docker-enabled applications across your cluster based on your resource needs and availability requirements.

For more information on ECS, please visit the link:

- <https://aws.amazon.com/ecs/details/>

NEW QUESTION 204

You are a DevOps Engineer and are designing an Opswork stack in AWS. The company has some custom recipes that are part of their on-premise Chef configuration. These same recipes need to be run whenever an instance is launched in Opsworks. Which of the following steps need to be carried out to ensure this requirement gets fulfilled. Choose 2 answers from the options given below

- A. Ensure the custom cookbooks option is set in Opswork stack.
- B. Ensure the custom cookbooks option is set in Opswork layer.
- C. Ensure the recipe is placed as part of the Setup Lifecycle event as part of the Layer setting.
- D. Ensure the recipe is placed as part of the Setup Lifecycle event as part of the Stack setting.

Answer: AC

Explanation:

The AWS Documentation mentions the below

Each layer has a set of built-in recipes assigned to each lifecycle event, although some layers lack Undeploy recipes. When a lifecycle event occurs on an instance,

AWS OpsWorks Stacks runs the appropriate set of recipes for the associated layer.

The below diagram shows a snapshot in the stack settings where you can enable custom cookbooks

For more information on automating recipe's, please visit the below URL: <http://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-assigningcustom.html>

NEW QUESTION 207

Which of the following file needs to be included along with your source code binaries when your application uses the EC2/On-Premises compute platform, and deploy it using the AWS Code Deploy service.

- A. appspec.yml
- B. appconfig.yml
- C. appspec.json
- D. appconfig.json

Answer: A

Explanation:

The AWS Documentation mentions the below

The application specification file (AppSpec file) is a YAML-formatted file used by AWS CodeDeploy to determine:

what it should install onto your instances from your application revision in Amazon S3 or GitHub. which lifecycle event hooks to run in response to deployment lifecycle events. An AppSpec file must be named appspec.yml and it must be placed in the root of an application's source code's directory structure. Otherwise, deployments will fail. For more information on the appspec file, please visit the below URL:

<http://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file.html>

Note: If you are deploying your code on AWS Lambda compute platform. An AppSpec file can be YAML-formatted or JSON-formatted. You can also enter the contents of an AppSpec file directly into AWS CodeDeploy console when you create a deployment.

However, this question is about along with your source code binaries on an EC2/On-Premises Compute Platform. So, an AppSpec file must be a YAML-formatted

file named appspec.yml and it must be placed in the root of the directory structure of an application's source code. Otherwise, deployments fail.

NEW QUESTION 208

When using EC2 instances with the Code Deploy service, which of the following are some of the pre- requisites to ensure that the EC2 instances can work with Code Deploy. Choose 2 answers from the options given below

- A. Ensurean 1AM role is attached to the instance so that it can work with the CodeDeploy Service.
- B. Ensurethe EC2 Instance is configured with Enhanced Networking
- C. Ensurethe EC2 Instance is placed in the default VPC
- D. Ensurethat the CodeDeploy agent is installed on the EC2 Instance

Answer: AD

Explanation:

This is mentioned in the AWS documentation

For more information on instances for CodeDeploy, please visit the below URL:

- <http://docs.aws.amazon.com/codedeploY/latest/userguide/instances.html>

NEW QUESTION 211

Your company has a number of Cloudformation stacks defined in AWS. As part of the routine housekeeping activity, a number of stacks have been targeted for deletion. But a few of the stacks are not getting deleted and are failing when you are trying to delete them. Which of the following could be valid reasons for this? Choose 2 answers from the options given below

- A. Thestacks were created with the wrong template versio
- B. Since the standardtemplate version is now higher, it is preventing the deletion of the stacks.You need to contact AWS support.
- C. Thestack has an S3 bucket defined which has objects present in it.
- D. Thestack has a EC2 Security Group which has EC2 Instances attached to it.
- E. Thestack consists of an EC2 resource which was created with a custom AMI.

Answer: BC

Explanation:

The AWS documentation mentions the below point

Some resources must be empty before they can be deleted. For example, you must delete all objects in an Amazon S3 bucket or remove all instances in an Amazon

CC2 security group before you can delete the bucket or security group

For more information on troubleshooting cloudformation stacks, please visit the below URL:

- <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/troubleshooting.html>

NEW QUESTION 212

You have a development team that is planning for continuous release cycles for their application. They want to use the AWS services available to be able to deploy a web application and also ensure they can rollback to previous versions fairly quickly. Which of the following options can be used to achieve this requirement. Choose 2 answers from the options given below

- A. Usethe Elastic beanstalk servic
- B. Use Application versions and upload therevisions of your applicatio
- C. Deploy the revisions accordingly and rollback toprior versions accordingly.
- D. Usethe Elastic beanstalk servic
- E. Create separate environments for eachapplication revisio
- F. Revert back to an environment incase the new environmentdoes not work.
- G. Usethe Opswork service to deploy the web instance
- H. Deploy the app to the Opsworkweb laye
- I. Rollback using the Deploy app in Opswor
- J. Usethe Cloudformation servic

K. Create separate templates for each application revision and deploy them accordingly.

Answer: AC

Explanation:

The AWS documentation mentions the following

In Elastic Beanstalk, an application version refers to a specific, labeled iteration of deployable code for a web application. An application version points to an Amazon

Simple Storage Service (Amazon S3) object that contains the deployable code such as a Java WAR file.

An application version is part of an application. Applications

can have many versions and each application version is unique. In a running environment, you can deploy any application version you already uploaded to the application or you can upload and immediately deploy a new application version. You might upload multiple application versions to test differences between one version of your web application and another.

For more information on Elastic beanstalk components, please refer to the below link:

? <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/concepts.components.html>

An AWS OpsWorks Stacks app represents code that you want to run on an application server. The code itself resides in a repository such as an Amazon S3 archive; the app contains the information required to deploy the code to the appropriate application server instances. For more information on Opswork apps, please refer to the below link:

• <http://docs.aws.amazon.com/opsworks/latest/userguide/workingapps.html>

Option B is incorrect. Our scenario is focusing on continuous development and continuous releases of the application versions. Since this is going to be an ongoing process, it is a best practice to upload the revision of your application and if required roll back to previous version.

Option D is incorrect. This question gives importance to the application hosted on the infrastructure. "They want to use the AWS services available to be able to deploy a web application and also ensure they can rollback to previous versions of the application quickly."

In this case, Cloud Formation provides a common language for you to describe and provision all the infrastructure resources in your cloud environment. Cloud Formation allows you to use a simple text file to model and provision, in an automated and secure manner, all the resources needed for your applications across all regions and accounts.

Hence, Cloud Formation is nothing to do with an application hosted on the infrastructure.

NEW QUESTION 217

Which of the following commands for the elastic beanstalk CLI can be used to create the current application into the specified environment?

- A. ebcreate
- B. ebstart
- C. enenv
- D. enapp

Answer: A

Explanation:

Differences from Version 3 of EB CLI

CB is a command line interface (CLI) tool for Elastic Beanstalk that you can use to deploy applications quickly and more easily. The latest version of CB was introduced by Elastic Beanstalk in CB CLI 3. Although Elastic Beanstalk still supports CB 2.6 for customers who previously installed and continue to use it, you should migrate to the latest version of CB CLI 3, as it can manage environments that you launched using CB CLI 2.6 or earlier versions of CB CLI. CB CLI automatically retrieves settings from an environment created using CB if the environment is running. Note that CB CLI 3 does not store option settings locally, as in earlier versions.

CB CLI introduces the commands `eb create`, `eb deploy`, `eb open`, `eb console`, `eb scale`, `eb setenv`, `eb config`, `eb terminate`, `eb clone`, `eb list`, `eb use`, `eb printenv`, and `eb ssh`. In CB CLI 3.1 or later, you can also use the `eb swap` command. In CB CLI 3.2 only, you can use the `eb abort`, `eb platform`, and `eb upgrade` commands. In addition to these new commands, CB CLI 3 commands differ from CB CLI 2.6 commands in several cases:

1. `eb init` - Use `eb init` to create an `.elasticbeanstalk` directory in an existing project directory and create a new Elastic Beanstalk application for the project. Unlike with previous versions, CB CLI 3 and later versions do not prompt you to create an environment.
2. `eb start` - CB CLI 3 does not include the command `eb start`. Use `eb create` to create an environment.
3. `eb stop` - CB CLI 3 does not include the command `eb stop`. Use `eb terminate` to completely terminate an environment and clean up.
4. `eb push` and `git aws.push` - CB CLI 3 does not include the commands `eb push` or `git aws.push`. Use `eb deploy` to update your application code.
5. `eb update` - CB CLI 3 does not include the command `eb update`. Use `eb config` to update an environment.
6. `eb branch` - CB CLI 3 does not include the command `eb branch`.

For more information about using CB CLI 3 commands to create and manage an application, see CB CLI Command Reference. For a command reference for CB 2.6, see CB CLI 2 Commands. For a walkthrough of how to deploy a sample application using CB CLI 3, see Managing Elastic Beanstalk environments with the CB CLI. For a walkthrough of how to deploy a sample application using `eb` 2.6, see Getting Started with `Cb`. For a walkthrough of how to use CB 2.6 to map a Git branch to a specific environment, see Deploying a Git Branch to a Specific environment. <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/eb-cli.html#eb-cli2-differences> Note: Additionally, CB CLI 2.6 has been deprecated. It has been replaced by AWS CLI <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/eb-cli3.html> We will replace this question soon.

NEW QUESTION 218

You are a Devops engineer for your company. You have been instructed to deploy docker containers using the Opswork service. How could you achieve this? Choose 2 answers from the options given below

- A. Use custom cookbooks for your Opswork stack and provide the Git repository which has the chef recipes for the Docker container
- B. ^
- C. Use Elastic beanstalk to deploy docker containers since this is not possible in Opswork
- D. Then attach the elastic beanstalk environment as a layer in Opswork.
- E. Use Cloudformation to deploy docker containers since this is not possible in Opswork
- F. Then attach the Cloudformation resources as a layer in Opswork.
- G. In the App for Opswork deployment, specify the git url for the recipes which will deploy the applications in the docker environment.

Answer: AD

Explanation:

This is mentioned in the AWS documentation

AWS OpsWorks lets you deploy and manage application of all shapes and sizes. Ops Works layers let you create blueprints for EC2 instances to install and configure any software that you want.

For more information on Opswork and Docker, please refer to the below link:

- <https://aws.amazon.com/blogs/devops/running-docker-on-aws-opsworks/>

NEW QUESTION 222

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DOP-C01 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DOP-C01 Product From:

<https://www.2passeasy.com/dumps/DOP-C01/>

Money Back Guarantee

DOP-C01 Practice Exam Features:

- * DOP-C01 Questions and Answers Updated Frequently
- * DOP-C01 Practice Questions Verified by Expert Senior Certified Staff
- * DOP-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * DOP-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year