

## Exam Questions DOP-C01

AWS Certified DevOps Engineer- Professional

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



### NEW QUESTION 1

You have an application which consists of EC2 instances in an Auto Scaling group. Between a particular time frame every day, there is an increase in traffic to your website. Hence users are complaining of a poor response time on the application. You have configured your Auto Scaling group to deploy one new EC2 instance when CPU utilization is greater than 60% for 2 consecutive periods of 5 minutes. What is the least cost-effective way to resolve this problem?

- A. Decrease the consecutive number of collection periods
- B. Increase the minimum number of instances in the Auto Scaling group
- C. Decrease the collection period to ten minutes
- D. Decrease the threshold CPU utilization percentage at which to deploy a new instance

**Answer: B**

#### Explanation:

If you increase the minimum number of instances, then they will be running even though the load is not high on the website. Hence you are incurring cost even though there is no need.

All of the remaining options are possible options which can be used to increase the number of instances on a high load.

For more information on On-demand scaling, please refer to the below link: <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scale-based-on-demand.html>

Note: The tricky part where the question is asking for 'least cost effective way'. You got the design consideration correctly but need to be careful on how the question is phrased.

### NEW QUESTION 2

You currently have the following setup in AWS

- 1) An Elastic Load Balancer
- 2) Auto Scaling Group which launches EC2 Instances
- 3) AMIs with your code pre-installed

You want to deploy the updates of your app to only a certain number of users. You want to have a cost-effective solution. You should also be able to revert back quickly. Which of the below solutions is the most feasible one?

- A. Create a second ELB, and a new Auto Scaling Group assigned a new Launch Configuratio
- B. Create a new AMI with the updated ap
- C. Use Route53 Weighted Round Robin records to adjust the proportion of traffic hitting the two ELBs.
- D. Create new AM Is with the new ap
- E. Then use the new EC2 instances in half proportion to the older instances.
- F. Redeploy with AWS Elastic Beanstalk and Elastic Beanstalk version
- G. Use Route 53 Weighted Round Robin records to adjust the proportion of traffic hitting the two ELBs
- H. Create a full second stack of instances, cut the DNS over to the new stack of instances, and change the DNS back if a rollback is needed.

**Answer: A**

#### Explanation:

The Weighted Routing policy of Route53 can be used to direct a proportion of traffic to your application. The best option is to create a second CLB, attach the new Autoscaling Group and then use Route53 to divert the traffic.

Option B is wrong because just having EC2 instances running with the new code will not help.

Option C is wrong because Clastic beanstalk is good for development environments, and also there is no mention of having 2 environments where environment url's can be swapped.

Option D is wrong because you still need Route53 to split the traffic.

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

### NEW QUESTION 3

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 4

You have just recently deployed an application on EC2 instances behind an ELB. After a couple of weeks, customers are complaining on receiving errors from the application. You want to diagnose the errors and are trying to get errors from the ELB access logs. But the ELB access logs are empty. What is the reason for this.

- A. You do not have the appropriate permissions to access the logs
- B. You do not have your CloudWatch metrics correctly configured
- C. ELB Access logs are only available for a maximum of one week.

D. Access logging is an optional feature of Elastic Load Balancing that is disabled by default

**Answer:** D

**Explanation:**

Elastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot issues.

Access logging is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logging for your load balancer, Elastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify. You can disable access logging at any time.

For more information on CLB access logs, please refer to the below document link: from AWS  
<http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/access-log-collection.html>

**NEW QUESTION 5**

You have deployed an application to AWS which makes use of Autoscaling to launch new instances. You now want to change the instance type for the new instances. Which of the following is one of the action items to achieve this deployment?

- A. Use Elastic Beanstalk to deploy the new application with the new instance type
- B. Use CloudFormation to deploy the new application with the new instance type
- C. Create a new launch configuration with the new instance type
- D. Create new EC2 instances with the new instance type and attach it to the Autoscaling Group

**Answer:** C

**Explanation:**

The ideal way is to create a new launch configuration, attach it to the existing Auto Scaling group, and terminate the running instances.

Option A is invalid because Elastic beanstalk cannot launch new instances on demand. Since the current scenario requires Autoscaling, this is not the ideal option

Option B is invalid because this will be a maintenance overhead, since you just have an Autoscaling Group. There is no need to create a whole CloudFormation template for this.

Option D is invalid because Autoscaling Group will still launch EC2 instances with the older launch configuration

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

**NEW QUESTION 6**

You have an ELB setup in AWS with EC2 instances running behind it. You have been requested to monitor the incoming connections to the ELB. Which of the below options can suffice this requirement?

- A. Use AWS CloudTrail with your load balancer
- B. Enable access logs on the load balancer
- C. Use a CloudWatch Logs Agent
- D. Create a custom metric CloudWatch filter on your load balancer

**Answer:** B

**Explanation:**

Elastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses.

You can use these access logs to analyze traffic patterns and to troubleshoot issues.

Option A is invalid because this service will monitor all AWS services Option C and D are invalid since CLB already provides a logging feature.

For more information on ELB access logs, please refer to the below document link: from AWS  
<http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/access-log-collection.html>

**NEW QUESTION 7**

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 send 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 send 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](http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US_SetupSNS.html)

#### NEW QUESTION 8

You are responsible for your company's large multi-tiered Windows-based web application running on Amazon EC2 instances situated behind a load balancer. While reviewing metrics, you've started noticing an upwards trend for slow customer page load time. Your manager has asked you to come up with a solution to ensure that customer load time is not affected by too many requests per second. Which technique would you use to solve this issue?

- A. Re-deploy your infrastructure usingan AWS CloudFormation templat
- B. Configure Elastic Load Balancing health checks to initiate a new AWS CloudFormation stack when health checks return failed.
- C. Re-deploy your infrastructure using an AWS CloudFormation templat
- D. Spin up a second AWS CloudFormation stac
- E. Configure Elastic Load Balancing SpillOver functionality to spill over any slow connections to the second AWS CloudFormation stack.
- F. Re-deploy your infrastructure using AWS CloudFormation, Elastic Beanstalk, and Auto Scalin
- G. Setup your Auto Scalinggroup policies to scale based on the number of requests per second as well as the current customer load tim
- H. •>/D- Re-deploy your application using an Auto Scaling templat
- I. Configure the Auto Scaling template to spin up a new Elastic Beanstalk application when the customer load time surpasses your threshold.

**Answer:** C

#### Explanation:

Auto Scaling helps you ensure that you have the correct number of Amazon CC2 instances available to handle the load for your application. You create collections of  
CC2 instances, called Auto Scaling groups. You can specify the minimum number of instances in each Auto Scaling group, and Auto Scaling ensures that your group  
never goes below this size. You can specify the maximum number of instances in each Auto Scaling group, and Auto Scaling ensures that yourgroup never goes above this size. If you specify the desired capacity, either when you create the group or at any time thereafter. Auto Scaling ensures that yourgroup has this many instances. If you specify scaling policies, then Auto Scaling can launch or terminate instances as demand on your application increases or decreases.  
Option A and B are invalid because Autoscaling is required to solve the issue to ensure the application can handle high traffic loads.  
Option D is invalid because there is no Autoscaling template.  
For more information on Autoscaling, please refer to the below document link: from AWS <http://docs.aws.amazon.com/autoscaling/latest/userguide/WhatIsAutoScaling.html>

#### NEW QUESTION 9

During metric analysis, your team has determined that the company's website during peak hours is experiencing response times higher than anticipated. You currently rely on Auto Scaling to make sure that you are scaling your environment during peak windows. How can you improve your Auto Scaling policy to reduce this high response time? Choose 2 answers.

- A. Push custom metrics to CloudWatch to monitor your CPU and network bandwidth from your servers, which will allow your Auto Scaling policy to have betterfine-grain insight.
- B. IncreaseyourAutoScalinggroup'snumberofmaxservers.
- C. Create a script that runs and monitors your servers; when it detects an anomaly in load, it posts to an Amazon SNS topic that triggers Elastic Load Balancing to add more servers to the load balancer.
- D. Push custom metrics to CloudWatch for your application that include more detailed information about your web application, such as how many requests it is handling and how many are waiting to be processed.

**Answer:** BD

#### Explanation:

Option B makes sense because maybe the max servers is low hence the application cannot handle the peak load.  
Option D helps in ensuring Autoscaling can scale the group on the right metrics.  
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 10

You currently run your infrastructure on Amazon EC2 instances behind an Auto Scalinggroup. All logs for your application are currentl\ written to ephemeral storage. Recently your company experienced a major bug in the code that made it through testing and was ultimately deployed to your fleet. This bug triggered your Auto Scalinggroup to scale up and back down before you could successfully retrieve the logs off your server to better assist you in troubleshooting the bug. Which technique should you use to make sure you are able to review your logs after your instances have shut down?

- A. Configure the ephemeral policies on your Auto Scaling group to back up on terminate.
- B. Configure your Auto Scaling policies to create a snapshot of all ephemeral storage on terminate.
- C. Install the CloudWatch Logs Agent on your AMI, and configure CloudWatch Logs Agent to stream your logs.V
- D. Install the CloudWatch monitoring agent on your AMI, and set up new SNS alert for CloudWatch metrics that triggers the CloudWatch monitoring agent to backup all logs on the ephemeral drive.

**Answer:** C

#### Explanation:

You can use Cloud Watch Logs to monitor applications and systems using log data. For example,  
CloudWatch Logs can track the number of errors that occur in your  
application logs and send you a notification whenever the rate of errors exceeds a threshold you specify. CloudWatch Logs uses your log data for monitoring; so,  
no  
code changes are required.  
Option A and B are invalid because Autoscaling policies are not designed for these purposes. Option D is invalid because you use Cloudwatch Logs Agent and not the monitoring agent. For more information on Cloudwatch logs, please refer to the below link:



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

#### NEW QUESTION 10

You have a code repository that uses Amazon S3 as a data store. During a recent audit of your security controls, some concerns were raised about maintaining the integrity of the data in the Amazon S3 bucket. Another concern was raised around securely deploying code from Amazon S3 to applications running on Amazon EC2 in a virtual private cloud. What are some measures that you can implement to mitigate these concerns? Choose two answers from the options given below.

- A. Add an Amazon S3 bucket policy with a condition statement to allow access only from Amazon EC2 instances with RFC 1918 IP addresses and enable bucket versioning.
- B. Add an Amazon S3 bucket policy with a condition statement that requires multi-factor authentication in order to delete objects and enable bucket versioning.
- C. Use a configuration management service to deploy AWS Identity and Access Management user credentials to the Amazon EC2 instance
- D. Use these credentials to securely access the Amazon S3 bucket when deploying code.
- E. Create an Amazon Identity and Access Management role with authorization to access the Amazon S3 bucket, and launch all of your application's Amazon EC2 instances with this role.
- F. Use AWS Data Pipeline to lifecycle the data in your Amazon S3 bucket to Amazon Glacier on a weekly basis.
- G. Use AWS Data Pipeline with multi-factor authentication to securely deploy code from the Amazon S3 bucket to your Amazon EC2 instances.

**Answer:** BD

#### Explanation:

You can add another layer of protection by enabling MFA Delete on a versioned bucket. Once you do so, you must provide your AWS account's access keys and a valid code from the account's MFA device in order to permanently delete an object version or suspend or reactivate versioning on the bucket. For more information on MFA please refer to the below link: <https://aws.amazon.com/blogs/security/securing-access-to-aws-using-mfa-part-3/> IAM roles are designed so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that the applications use. Instead of creating and distributing your AWS credentials, you can delegate permission to make API requests using IAM roles For more information on Roles for EC2 please refer to the below link: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html> Option A is invalid because this will not address either the integrity or security concern completely. Option C is invalid because user credentials should never be used in EC2 instances to access AWS resources. Option E and F are invalid because AWS Pipeline is an unnecessary overhead when you already have inbuilt controls to manage security for S3.

#### NEW QUESTION 14

The operations team and the development team want a single place to view both operating system and application logs. How should you implement this using AWS services? Choose two from the options below

- A. Using AWS CloudFormation, create a Cloud Watch Logs LogGroup and send the operating system and application logs of interest using the Cloud Watch Logs Agent.
- B. Using AWS CloudFormation and configuration management, set up remote logging to send events via UDP packets to CloudTrail.
- C. Using configuration management, set up remote logging to send events to Amazon Kinesis and insert these into Amazon CloudSearch or Amazon Redshift, depending on available analytic tools.
- D. Using AWS CloudFormation, merge the application logs with the operating system logs, and use IAM Roles to allow both teams to have access to view console output from Amazon EC2.

**Answer:** AC

#### Explanation:

Option B is invalid because Cloudtrail is not designed specifically to take in UDP packets  
Option D is invalid because there are already Cloudwatch logs available, so there is no need to have specific logs designed for this.  
You can use Amazon CloudWatch Logs to monitor, store, and access your log files from Amazon Elastic Compute Cloud (Amazon EC2) instances, AWS CloudTrail, and other sources. You can then retrieve the associated log data from CloudWatch Logs. For more information on Cloudwatch logs please refer to the below link: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html> You can use Kinesis to process those logs  
For more information on Amazon Kinesis please refer to the below link: <http://docs.aws.amazon.com/streams/latest/dev/introduction.html>

#### NEW QUESTION 19

You have the following application to be setup in AWS

- 1) A web tier hosted on EC2 Instances
- 2) Session data to be written to DynamoDB
- 3) Log files to be written to Microsoft SQL Server

How can you allow an application to write data to a DynamoDB table?

- A. Add an IAM user to a running EC2 instance.
- B. Add an IAM user that allows write access to the DynamoDB table.
- C. Create an IAM role that allows read access to the DynamoDB table.
- D. Create an IAM role that allows write access to the DynamoDB table.

**Answer:** D

#### Explanation:

IAM roles are designed so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that the applications use. Instead of creating and distributing your AWS credentials For more information on IAM Roles please refer to the below link: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

#### NEW QUESTION 20

You are a DevOps engineer for a company. You have been requested to create a rolling deployment solution that is cost-effective with minimal downtime. How should you achieve this? Choose two answers from the options below

- A. Re-deploy your application using a CloudFormation template to deploy Elastic Beanstalk
- B. Re-deploy with a CloudFormation template, define update policies on Auto Scaling groups in your CloudFormation template
- C. Use UpdatePolicy attribute to specify how CloudFormation handles updates to Auto Scaling Group resource.
- D. After each stack is deployed, tear down the old stack

**Answer:** BC

**Explanation:**

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 it is not efficient to use Cloudformation to use Elastic Beanstalk. Option D is invalid because this is an inefficient process to tear down stacks when there are stack policies available. 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 22**

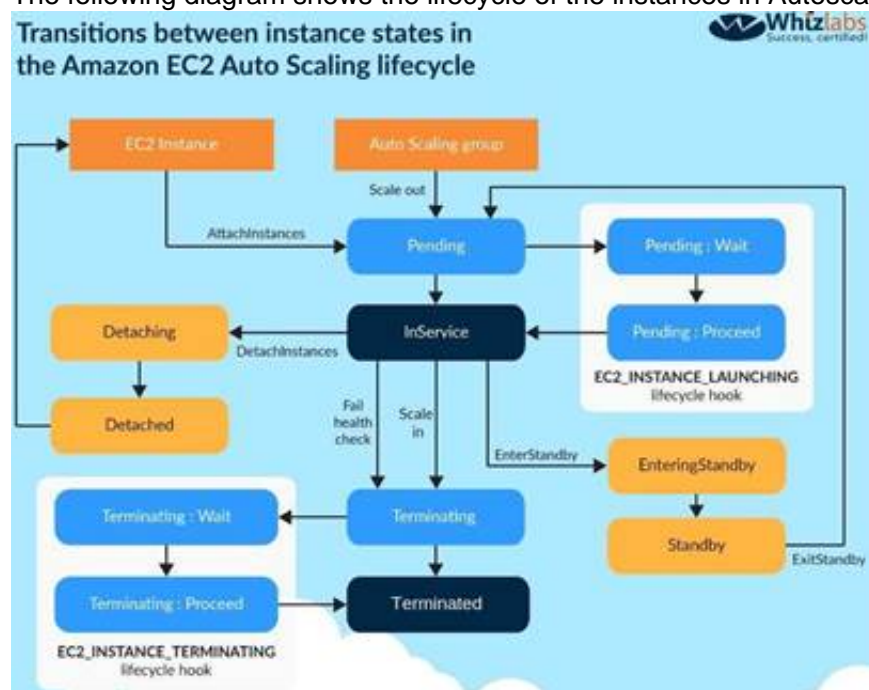
You have an Auto Scaling group of Instances that processes messages from an Amazon Simple Queue Service (SQS) queue. The group scales on the size of the queue. Processing Involves calling a third- party web service. The web service is complaining about the number of failed and repeated calls it is receiving 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 25**

If your application performs operations or workflows that take a long time to complete, what service can the Elastic Beanstalk environment do for you?

- A. Manages a Amazon SQS queue and running a daemon process on each instance
- B. Manages a Amazon SNS Topic and running a daemon process on each instance
- C. Manages Lambda functions and running a daemon process on each instance
- D. Manages the ELB and running a daemon process on each instance

**Answer:** A

**Explanation:**

Elastic Beanstalk simplifies this process by managing the Amazon SQS queue and running a daemon process on each instance that reads from the queue for you. When the daemon pulls an item from the queue, it sends an HTTP POST request locally to <http://localhost/> with the contents of the queue message in the body. All that your application needs to do is perform the long-running task in response to the POST.

For more information Elastic Beanstalk managing worker environments, please visit the below URL:

? <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.htm> |

**NEW QUESTION 26**

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 scrip
- 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 31**

You have a multi-docker environment that you want to deploy to AWS. Which of the following configuration files can be used to deploy a set of Docker containers as an Elastic Beanstalk application?

- A. `Dockerrun.awsjson`
- B. `.ebextensions`
- C. `Dockerrun.json`
- D. `Dockerfile`

**Answer:** A

**Explanation:**

A `Dockerrun.aws.json` file is an Elastic Beanstalk-specific JSON file that describes how to deploy a set of Docker containers as an Elastic Beanstalk application.

You can use a `Dockerrun.aws.json` file for a multicontainer Docker environment.

`Dockerrun.aws.json` describes the containers to deploy to each container instance in the environment as well as the data volumes to create on the host instance for

the containers to mount. [http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create\\_deploy\\_docker\\_v2config.html](http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_v2config.html)

**NEW QUESTION 33**

Your current log analysis application takes more than four hours to generate a report of the top 10 users of your web application. You have been asked to implement a system that can report this information in real time, ensure that the report is always up to date, and handle increases in the number of requests to your web application. Choose the option that is cost-effective and can fulfill the requirements.

- A. Publish your data to Cloud Watch Logs, and configure your application to autoscale to handle the load on demand.
- B. Publish your log data to an Amazon S3 bucket
- C. Use AWS CloudFormation to create an Auto Scaling group to scale your post-processing application which is configured to pull down your log files stored on Amazon S3.
- D. Post your log data to an Amazon Kinesis data stream, and subscribe your log-processing application so that is configured to process your logging data.
- E. Create a multi-AZ Amazon RDS MySQL cluster, post the logging data to MySQL, and run a map reduce job to retrieve the required information on user counts.

**Answer:** C

**Explanation:**

When you see Amazon Kinesis as an option, this becomes the ideal option to process data in real time.

Amazon Kinesis makes it easy to collect, process, and analyze real-time, streaming data so you can get timely insights and react quickly to new information.

Amazon

Kinesis offers key capabilities to cost effectively process streaming data at any scale, along with the flexibility to choose the tools that best suit the requirements of your application. With Amazon Kinesis, you can ingest real-time data such as application logs, website clickstreams, IoT telemetry data, and more into your databases, data lakes and data warehouses, or build your own real-time applications using this data. For more information on Amazon Kinesis, please visit the below URL:

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

**NEW QUESTION 37**

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 templat
- 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 39**

After reviewing the last quarter's monthly bills, management has noticed an increase in the overall bill from Amazon. After researching this increase in cost, you discovered that one of your new services is doing a lot of GET Bucket API calls to Amazon S3 to build a metadata cache of all objects in the applications bucket. Your boss has asked you to come up with a new cost-effective way to help reduce the amount of these new GET Bucket API calls. What process should you use to help mitigate the cost?

- A. Update your Amazon S3 buckets' lifecycle policies to automatically push a list of objects to a new bucket, and use this list to view objects associated with the application's bucket.
- B. Create a new DynamoDB tabl
- C. Use the new DynamoDB table to store all metadata about all objects uploaded to Amazon S3. Any time a new object is uploaded, update the application's internal Amazon S3 object metadata cache from DynamoDB.C Using Amazon SNS, create a notification on any new Amazon S3 objects that automatical ly updates a new DynamoDB table to store all metadata about the new objec
- D. Subscribe the application to the Amazon SNS topic to update its internal Amazon S3 object metadata cache from the DynamoDB tabl
- E. ^/
- F. Upload all files to an ElastiCache file cache serve
- G. Update your application to now read all file metadata from the ElastiCache file cache server, and configure the ElastiCache policies to push all files to Amazon S3 for long-term storage.

**Answer:** C

**Explanation:**

Option A is an invalid option since Lifecycle policies are normally used for expiration of objects or archival of objects.

Option B is partially correct where you store the data in DynamoDB, but then the number of GET requests would still be high if the entire DynamoDB table had to be

traversed and each object compared and updated in S3.

Option D is invalid because uploading all files to Clastic Cache is not an ideal solution.

The best option is to have a notification which can then trigger an update to the application to update the DynamoDB table accordingly.

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

- ? <https://aws.amazon.com/blogs/compute/619/>

**NEW QUESTION 42**

As part of your continuous deployment process, your application undergoes an I/O load performance test before it is deployed to production using new AMIs. The application uses one Amazon Elastic Block Store (EBS) PIOPS volume per instance and requires consistent I/O performance. Which of the following must be carried out to ensure that I/O load performance tests yield the correct results in a repeatable manner?

- A. Ensure that the I/O block sizes for the test are randomly selected.
- B. Ensure that the Amazon EBS volumes have been pre-warmed by reading all the blocks before the test.
- C. Ensure that snapshots of the Amazon EBS volumes are created as a backup.
- D. Ensure that the Amazon EBS volume is encrypted.

**Answer:** B

**Explanation:**

During the AMI-creation process, Amazon CC2 creates snapshots of your instance's root volume and any other CBS volumes attached to your instance

New CBS volumes receive their maximum performance the moment that they are available and do not require initialization (formerly known as pre-warming).

However, storage blocks on volumes that were restored from snapshots must to initialized (pulled

down from Amazon S3 and written to the volume) before you can access the block. This preliminary action takes time and can cause a significant increase in the latency of an I/O operation the first time each block is accessed. For most applications, amortizing this cost over the lifetime of the volume is acceptable.

Option A is invalid because block sizes are predetermined and should not be randomly selected. Option C is invalid because this is part of continuous integration and hence volumes can be destroyed after the test and hence there should not be snapshots created unnecessarily

Option D is invalid because the encryption is a security feature and not part of load tests normally. For more information on CBS initialization please refer to the below link:

- <http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/ebs-initialize.html>

**NEW QUESTION 45**

You have a complex system that involves networking, IAM policies, and multiple, three-tier applications. You are still receiving requirements for the new system, so you don't yet know how many AWS components will be present in the final design. You want to start using AWS CloudFormation to define these AWS resources so that you can automate and version-control your infrastructure. How would you use AWS CloudFormation to provide agile new environments for your customers in a cost-effective, reliable manner?

- A. Manually create one template to encompass all the resources that you need for the system, so you only have a single template to version-control.
- B. Create multiple separate templates for each logical part of the system, create nested stacks in AWS CloudFormation, and maintain several templates to version-control



- C. ➤/
- D. Create multiple separate templates for each logical part of the system, and provide the outputs from one to the next using an Amazon Elastic Compute Cloud (EC2) instance running the SDK for finer granularity of control.
- E. Manually construct the networking layer using Amazon Virtual Private Cloud (VPC) because this does not change often, and then use AWS CloudFormation to define all other ephemeral resources.

**Answer:** B

**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::StackResource` in your template to reference other templates.

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

**NEW QUESTION 47**

You use Amazon Cloud Watch as your primary monitoring system for your web application. After a recent software deployment, your users are getting Intermittent 500 Internal Server Errors when using the web application. You want to create a Cloud Watch alarm, and notify an on-call engineer when these occur. How can you accomplish this using AWS services? Choose three answers from the options given below

- A. Deploy your web application as an AWS Elastic Beanstalk applicatio
- B. Use the default Elastic Beanstalk Cloudwatch metrics to capture 500 Internal Server Error
- C. Set a CloudWatch alarm on that metric.
- D. Install a CloudWatch Logs Agent on your servers to stream web application logs to CloudWatch.
- E. Use Amazon Simple Email Service to notify an on-call engineer when a CloudWatch alarm is triggered.
- F. Create a CloudWatch Logs group and define metric filters that capture 500 Internal Server Error
- G. Set a CloudWatch alarm on that metric.
- H. Use Amazon Simple Notification Service to notify an on-call engineer when a CloudWatch alarm is triggered.

**Answer:** BDE

**Explanation:**

You can use Cloud Watch Logs to monitor applications and systems using log data

Cloud Watch Logs uses your log data for monitoring; so, no code changes are required. For example, you can monitor application logs for specific literal terms (such as "NullPointerException") or count the number of occurrences of a literal term at a particular position in log data (such as "404" status codes in an Apache access log). When the term you are searching for is found, Cloud Watch Logs reports the data to a CloudWatch metric that you specify. Log data is encrypted while in transit and while it is at rest

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

Amazon CloudWatch 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 SNS and Cloudwatch logs please refer to the below link:

[http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US\\_SetupSNS.html](http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US_SetupSNS.html)

**NEW QUESTION 50**

You are using CloudFormation to launch an EC2 instance and then configure an application after the instance is launched. You need the stack creation of the ELB and Auto Scaling to wait until the EC2 instance is launched and configured properly. How do you do this?

- A. It is not possible for the stack creation to wait until one service is created and launched
- B. Use the WaitCondition resource to hold the creation of the other dependent resources
- C. Use a CreationPolicy to wait for the creation of the other dependent resources ➤/
- D. Use the HoldCondition resource to hold the creation of the other dependent resources

**Answer:** C

**Explanation:**

When you provision an Amazon EC2 instance in an AWS Cloud Formation 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 Creation Policy 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.

A Creation Policy instructs CloudFormation to wait on an instance until CloudFormation receives the specified number of signals

Option A is invalid because this is possible

Option B is invalid because this is used to make AWS CloudFormation pause the creation of a stack and wait for a signal before it continues to create the stack

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

- <https://aws.amazon.com/blogs/devops/use-a-creationpolicy-to-wait-for-on-instance-configurations/>

**NEW QUESTION 53**

Your development team wants account-level access to production instances in order to do live debugging of a highly secure environment. Which of the following should you do?

- A. Place the credentials provided by Amazon Elastic Compute Cloud (EC2) into a secure Amazon Simple Storage Service (S3) bucket with encryption enable
- B. Assign AWS Identity and Access Management (IAM) users to each developer so they can download the credentials file.
- C. Place an internally created private key into a secure S3 bucket with server-side encryption using customer keys and configuration management, create a service account on all the instances using this private key, and assign IAM users to each developer so they can download the file.
- D. Place each developer's own public key into a private S3 bucket, use instance profiles and configuration management to create a user account for each developer on all instances, and place the user's public keys into the appropriate account
- E. ➤/
- F. Place the credentials provided by Amazon EC2 onto an MFA encrypted USB drive, and physically share it with each developer so that the private key never leaves the office.

**Answer:** C

**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.

A private S3 bucket can be created for each developer, the keys can be stored in the bucket and then assigned to the instance profile.

Option A and D are invalid, because the credentials should not be provided by a AWS EC2 Instance. Option B is invalid because you would not create a service account, instead you should create an instance profile.

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

• [http://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles\\_use\\_switch-role-ec2-instance-profiles.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2-instance-profiles.html)

**NEW QUESTION 57**

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 61**

As an architect you have decided to use CloudFormation instead of OpsWorks or Elastic Beanstalk for deploying the applications in your company. Unfortunately, you have discovered that there is a resource type that is not supported by CloudFormation. What can you do to get around this.

- A. Specify more mappings and separate your template into multiple templates by using nested stacks.
- B. Create a custom resource type using template developer, custom resource template, and CloudFormation.
- C. \*/
- D. Specify the custom resource by separating your template into multiple templates by using nested stacks.
- E. Use a configuration management tool such as Chef, Puppet, or Ansible.

**Answer:** B

**Explanation:**

Custom resources enable you to write custom provisioning logic in templates that AWS CloudFormation runs anytime you create, update (if you changed the custom resource), or delete stacks. For example, you might want to include resources that aren't available as AWS CloudFormation resource types. You can include those resources by using custom resources. That way you can still manage all your related resources in a single stack.

For more information on custom resources in CloudFormation please visit the below URL:

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

**NEW QUESTION 63**

You work for an insurance company and are responsible for the day-to-day operations of your company's online quote system used to provide insurance quotes to members of the public. Your company wants to use the application logs generated by the system to better understand customer behavior. Industry regulations also require that you retain all application logs for the system indefinitely in order to investigate fraudulent claims in the future. You have been tasked with designing a log management system with the following requirements:

- All log entries must be retained by the system, even during unplanned instance failure.
- The customer insight team requires immediate access to the logs from the past seven days.
- The fraud investigation team requires access to all historic logs, but will wait up to 24 hours before these logs are available.

How would you meet these requirements in a cost-effective manner? Choose three answers from the options below

- A. Configure your application to write logs to the instance's ephemeral disk, because this storage is free and has good write performance.
- B. Create a script that moves the logs from the instance to Amazon S3 once an hour.
- C. Write a script that is configured to be executed when the instance is stopped or terminated and that will upload any remaining logs on the instance to Amazon S3.
- D. Create an Amazon S3 lifecycle configuration to move log files from Amazon S3 to Amazon Glacier after seven days.
- E. Configure your application to write logs to the instance's default Amazon EBS boot volume, because this storage already exists.
- F. Create a script that moves the logs from the instance to Amazon S3 once an hour.
- G. Configure your application to write logs to a separate Amazon EBS volume with the "delete on termination" field set to false.
- H. Create a script that moves the logs from the instance to Amazon S3 once an hour.
- I. Create a housekeeping script that runs on a T2 micro instance managed by an Auto Scaling group for high availability.
- J. The script uses the AWS API to identify any unattached Amazon EBS volumes containing log files.
- K. Your housekeeping script will mount the Amazon EBS volume, upload all logs to Amazon S3, and then delete the volume.

**Answer:** CEF

**Explanation:**

Since all logs need to be stored indefinitely, Glacier is the best option for this. One can use Lifecycle events to stream the data from S3 to Glacier.

Lifecycle configuration enables you to specify the lifecycle management of objects in a bucket. The configuration is a set of one or more rules, where each rule defines an action for Amazon S3 to apply to a group of objects. These actions can be classified as

follows:

- Transition actions - In which you define when objects transition to another storage class. For example, you may choose to transition objects to the STANDARDIA (infrequent access) storage class 30 days after creation, or archive objects to the GLACIER storage class one year after creation.
- Expiration actions - In which you specify when the objects expire. Then Amazon S3 deletes the expired objects on your behalf. For more information on Lifecycle events, please refer to the below link:
- <http://docs.aws.amazon.com/AmazonS3/latest/dev/object-lifecycle-mgmt.html> | You can use scripts to put the logs onto a new volume and then transfer those logs to S3.

Note:

Moving the logs from CBS volume to S3 we have some custom scripts running in the background. In order to ensure the minimum memory requirements for the OS and the applications for the script to execute we can use a cost effective ec2 instance.

Considering the computing resource requirements of the instance and the cost factor a t2micro instance can be used in this case.

The following link provides more information on various t2 instances. <https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/t2-instances.html>

Question is "How would you meet these requirements in a cost-effective manner? Choose three answers from the options below"

So here user has to choose the 3 options so that the requirement is fulfilled. So in the given 6 options, options C, C and F fulfill the requirement.

"The CC2s use CBS volumes and the logs are stored on CBS volumes those are marked for non-termination" - is one of the ways to fulfill requirement. So this shouldn't be an issue.

#### NEW QUESTION 64

You need to implement Blue/Green Deployment for several multi-tier web applications. Each of them has its individual infrastructure:

Amazon Elastic Compute Cloud (EC2) front-end servers, Amazon ElastiCache clusters, Amazon Simple Queue Service (SQS) queues, and Amazon Relational Database (RDS) Instances.


Which combination of services would give you the ability to control traffic between different deployed versions of your application?

- Create one AWS Elastic Beanstalk application and all AWS resources (using configuration files inside the application source bundle) for each web application.
- New versions would be deployed using Elastic Beanstalk environments and using the Swap URLs feature.
- Using AWS CloudFormation templates, create one Elastic Beanstalk application and all AWS resources (in the same template) for each web application.
- New versions would be deployed using AWS CloudFormation templates to create new Elastic Beanstalk environments, and traffic would be balanced between them using weighted Round Robin (WRR) records in Amazon Route 53.
- Using AWS CloudFormation templates, create one Elastic Beanstalk application and all AWS resources (in the same template) for each web application.
- New versions would be deployed updating a parameter on the CloudFormation template and passing it to the cfn-hup helper daemon, and traffic would be balanced between them using Weighted Round Robin (WRR) records in Amazon Route 53.
- Create one Elastic Beanstalk application and all AWS resources (using configuration files inside the application source bundle) for each web application.
- New versions would be deployed updating the Elastic Beanstalk application version for the current Elastic Beanstalk environment.

**Answer: B**

#### Explanation:

This is an example of Blue green deployment



With Amazon Route 53, you can define a percentage of traffic to go to the green environment and gradually update the weights until the green environment carries the full production traffic. A weighted distribution provides the ability to perform canary analysis where a small percentage of production traffic is introduced to a new environment. You can test the new code and monitor for errors, limiting the blast radius if any issues are encountered. It also allows the green environment to scale out to support the full production load if you're using Elastic Load Balancing.

When it's time to promote the green environment/stack into production, update DNS records to point to the green environment/stack's load balancer. You can also do this DNS flip gradually by using the Amazon Route 53 weighted routing policy. For more information on Blue green deployment, please refer to the link:

- [https://dOawsstatic.com/whitepapers/AWS\\_Blue\\_Green\\_Deployments.pdf](https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf)

#### NEW QUESTION 69

You have deployed an Elastic Beanstalk application in a new environment and want to save the current state of your environment in a document. You want to be able to restore your environment to the current state later or possibly create a new environment. You also want to make sure you have a restore point. How can you achieve this?

- Use CloudFormation templates
- Configuration Management Templates
- Saved Configurations
- Saved Templates

**Answer: C**

#### Explanation:

You can save your environment's configuration as an object in Amazon S3 that can be applied to other environments during environment creation, or applied to a running environment. Saved configurations are YAML formatted templates that define an environment's platform configuration, tier, configuration option settings, and tags.

For more information on Saved Configurations please refer to the below link:



- <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environment-configuration-savedconfig.html>

#### NEW QUESTION 71

You currently have an Auto Scaling group with an Elastic Load Balancer and need to phase out all instances and replace with a new instance type. What are 2 ways in which this can be achieved.

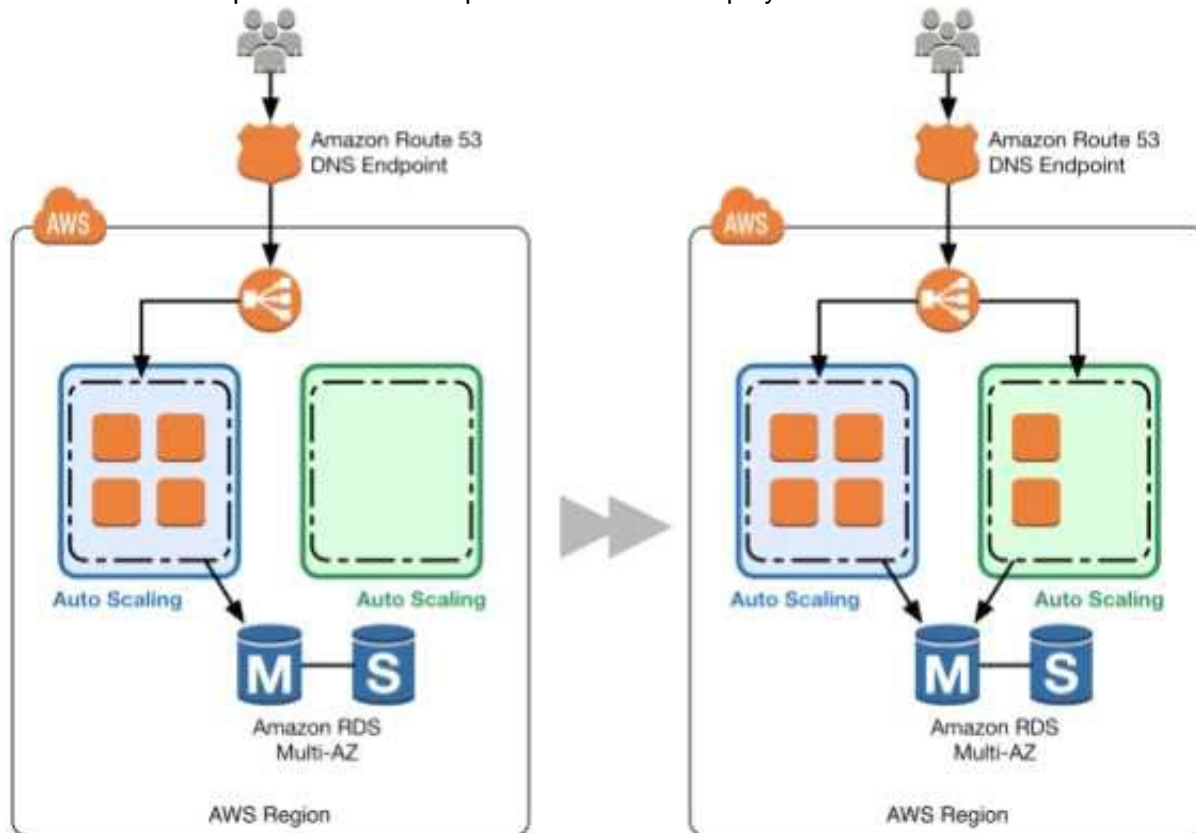
- A. Use Newest Instance to phase out all instances that use the previous configuration.
- B. Attach an additional ELB to your Auto Scaling configuration and phase in newer instances while removing older instances.
- C. Use OldestLaunchConfiguration to phase out all instances that use the previous configuration
- D. V
- E. Attach an additional Auto Scaling configuration behind the ELB and phase in newer instances while removing older instances.

**Answer: CD**

#### Explanation:

When using the OldestLaunchConfiguration policy Auto Scaling terminates instances that have the oldest launch configuration. This policy is useful when you're updating a group and phasing out the instances from a previous configuration.

For more information on Autoscaling instance termination, please visit the below URL: <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-instance-termination.html> Option D is an example of Blue Green Deployments.



A blue group carries the production load while a green group is staged and deployed with the new code. When it's time to deploy, you simply attach the green group to the existing load balancer to introduce traffic to the new environment. For HTTP/HTTPS listeners, the load balancer favors the green Auto Scaling group because it uses a least outstanding requests routing algorithm.

As you scale up the green Auto Scaling group, you can take blue Auto Scaling group instances out of service by either terminating them or putting them in Standby state.

For more information on Blue Green Deployments, please refer to the below document link: from AWS

- [https://dOawsstatic.com/whitepapers/AWS\\_Blue\\_Green\\_Deployments.pdf](https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf)

#### NEW QUESTION 75

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 hosting 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.wsamazon.com/IAM/latest/UserGuide/id\\_roles\\_providers\\_oidc.html](http://docs.wsamazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html)

#### NEW QUESTION 80

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

#### NEW QUESTION 84

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

#### NEW QUESTION 85

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 87**

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 91**

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

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

**Answer:** C

**Explanation:**

The best way is to distribute the instances across multiple AZ's to get the best and avoid a disaster scenario. With this scenario, you will always have a minimum of more than 8 servers even if one AZ were to go down. Even though A and D are also valid options, the best option when it comes to distribution is Option C. For more information on High Availability and Fault tolerance, please refer to the below link:

[https://media.amazonwebservices.com/architecturecenter/AWS\\_ac\\_ra\\_ftha\\_04.pdf](https://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_ftha_04.pdf)

**NEW QUESTION 94**

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 96**

You have an application hosted in AWS. You wanted to ensure that when certain thresholds are reached, a DevOps Engineer is notified. Choose 3 answers from the options given below

- A. Use CloudWatch Logs agent to send log data from the app to CloudWatch Logs from Amazon EC2 instances
- B. Pipe data from EC2 to the application logs using AWS Data Pipeline and CloudWatch
- C. Once a CloudWatch alarm is triggered, use SNS to notify the Senior DevOps Engineer.



D. Set the threshold your application can tolerate in a CloudWatch Logs group and link a CloudWatch alarm on that threshold.

**Answer:** ACD

**Explanation:**

You can use Cloud Watch Logs to monitor applications and systems using log data. For example, CloudWatch Logs can track the number of errors that occur in your application logs and send you a notification whenever the rate of errors exceeds a threshold you specify. CloudWatch Logs uses your log data for monitoring; so, no code changes are required. For example, you can monitor application logs for specific literal terms (such as "NullPointerException") or count the number of occurrences of a literal term at a particular position in log data (such as "404" status codes in an Apache access log). When the term you are searching for is found, CloudWatch Logs reports the data to a CloudWatch metric that you specify. For more information on Cloudwatch Logs please refer to the below link:  
<http://docs.ws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html>  
Amazon CloudWatch 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 Cloudwatch and SNS please refer to the below link:  
[http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US\\_SetupSNS.html](http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US_SetupSNS.html)

**NEW QUESTION 101**

Your company releases new features with high frequency while demanding high application availability. As part of the application's A/B testing, logs from each updated Amazon EC2 instance of the application need to be analyzed in near real-time, to ensure that the application is working flawlessly after each deployment. If the logs show any anomalous behavior, then the application version of the instance is changed to a more stable one. Which of the following methods should you use for shipping and analyzing the logs in a highly available manner?

- A. Ship the logs to Amazon S3 for durability and use Amazon EMR to analyze the logs in a batch manner each hour.
- B. Ship the logs to Amazon CloudWatch Logs and use Amazon EMR to analyze the logs in a batch manner each hour.
- C. Ship the logs to an Amazon Kinesis stream and have the consumers analyze the logs in a live manner.
- D. Ship the logs to a large Amazon EC2 instance and analyze the logs in a live manner.

**Answer:** C

**Explanation:**

Answer - C

You can use Kinesis Streams 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. Because the response time for the data intake and processing is in real time, the processing is typically lightweight.

The following are typical scenarios for using Kinesis Streams:

- Accelerated log and data feed intake and processing - You can have producers push data directly into a stream. For example, push system and application logs and they'll be available for processing in seconds. This prevents the log data from being lost if the front end or application server fails. Kinesis Streams provides accelerated data feed intake because you don't batch the data on the servers before you submit it for intake.
  - Real-time metrics and reporting - You can use data collected into Kinesis Streams for simple data analysis and reporting in real time. For example, your data-processing application can work on metrics and reporting for system and application logs as the data is streaming in, rather than wait to receive batches of data.
- For more information on Amazon Kinesis and SNS please refer to the below link:  
• <http://docs.aws.amazon.com/streams/latest/dev/introduction.html>

**NEW QUESTION 105**

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 106**

You are creating a new API for video game scores. Reads are 100 times more common than writes, and the top 1% of scores are read 100 times more frequently than the rest of the scores. What's the best design for this system, using DynamoDB?

- A. DynamoDB table with 100x higher read than write throughput, with CloudFront caching.
- B. DynamoDB table with roughly equal read and write throughput, with CloudFront caching.
- C. DynamoDB table with 100x higher read than write throughput, with ElastiCache caching.
- D. DynamoDB table with roughly equal read and write throughput, with ElastiCache caching.

**Answer:** D

**Explanation:**

Because the 100x read ratio is mostly driven by a small subset, with caching, only a roughly equal number of reads to writes will miss the cache, since the supermajority will hit the top 1% scores. Knowing we need to set the values roughly equal when using caching, we select AWS ElastiCache, because CloudFront cannot directly cache DynamoDB queries, and ElastiCache is an excellent in-memory cache for database queries, rather than a distributed proxy cache for content delivery.

For more information on DynamoDB table guidelines please refer to the below link:

- <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GuidelinesForTables.html>

#### NEW QUESTION 107

You are planning on using the Amazon RDS facility for Fault tolerance for your application. How does Amazon RDS Multi Availability Zone model work

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

**Answer: A**

#### Explanation:

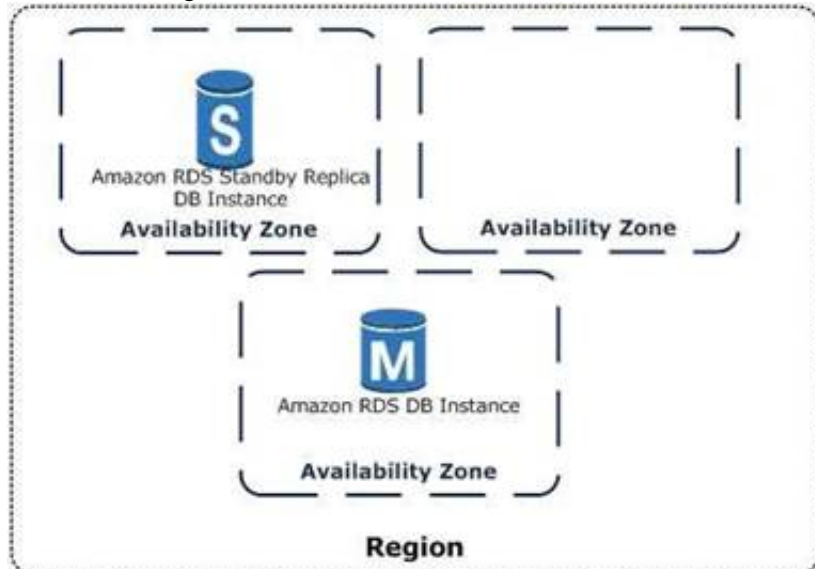
Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB)

Instances, making them a natural fit for production database

workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable.

In case of an infrastructure failure, Amazon RDS performs an automatic failover to the standby (or to a read replica in the case of Amazon Aurora), so that you can resume database operations as soon as the failover is complete.

The below diagram from the AWS documentation shows how this is configured



Option B is invalid because the replication is synchronous.

Option C and D are invalid because this is built around AZ and not regions. For more information on Multi-AZ RDS, please visit the below URL:

<https://aws.amazon.com/rds/details/multi-az/>

#### NEW QUESTION 109

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 Events 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 Events becomes aware of operational changes as they occur. Cloud Watch Events 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 110

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

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

**Answer: C**

#### Explanation:

The most likely cause is the fact you have hit the maximum of 5 Elastic IP's per region.

By default, all AWS accounts are limited to 5 Elastic IP addresses per region, because public (IPv4) Internet addresses are a scarce public resource. We strongly encourage you to use an Elastic IP address primarily for the ability to remap the address to another instance in the case of instance failure, and to use DNS hostnames for all other inter-node communication.

Option A is invalid because a AMI does not have a Development version tag. Option B is invalid because there is no flag for an EC2 Instance

Option D is invalid because there is a limit of 5 VPCs per region. For more information on Elastic IP's, please visit the below URL:  
• <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html>

#### NEW QUESTION 112

You are designing a service that aggregates clickstream data in batch and delivers reports to subscribers via email only once per week. Data is extremely spikey, geographically distributed, high- scale, and unpredictable. How should you design this system?

- A. Use a large RedShift cluster to perform the analysis, and a fleet of Lambdas to perform record inserts into the RedShift table
- B. Lambda will scale rapidly enough for the traffic spikes.
- C. Use a CloudFront distribution with access log delivery to S3. Clicks should be recorded as querystring GETs to the distribution
- D. Reports are built and sent by periodically running EMR jobs over the access logs in S3. C Use API Gateway invoking Lambdas which PutRecords into Kinesis, and EMR running Spark performing GetRecords on Kinesis to scale with spike
- E. Spark on EMR outputs the analysis to S3, which are sent out via email. D- Use AWS Elasticsearch service and EC2 Auto Scaling group
- F. The Autoscaling groups scale based on click throughput and stream into the Elasticsearch domain, which is also scalable
- G. Use Kibana to generate reports periodically.

**Answer: B**

#### Explanation:

When you look at building reports or analyzing data from a large data set, you need to consider CMR because this service is built on the Hadoop framework which is used to process large data sets.

The ideal approach to getting data onto CMR is to use S3. Since the Data is extremely spikey and geographically distributed, using edge locations via Cloudfront distributions is the best way to fetch the data.

Option A is invalid because RedShift is more of a petabyte storage cluster.

Option C is invalid because having both Kinesis and CMR for the job analysis is redundant. Option D is invalid because Elastic Search is not an option for processing records.

For more information on Amazon CMR, please visit the below URL:

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

#### NEW QUESTION 117

You want to pass queue messages that are 1GB each. How should you achieve 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. Use SQS's support for message partitioning and multi-part uploads on Amazon S3.
- E. Use AWS EFS as a shared pool storage medium
- F. Store filesystem pointers to the files on disk in the SQS message bodies.

**Answer: B**

#### Explanation:

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 processing large messages for SQS, please visit the below URL:

<http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-s3-messages.html>

#### NEW QUESTION 120

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

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

**Answer: A**

#### Explanation:

Amazon Elasticsearch Service makes it easy to deploy, operate, and scale Elasticsearch for log analytics, full text search, application monitoring, and more.

Amazon

Elasticsearch Service is a fully managed service that delivers Elasticsearch's easy-to-use APIs and real-time capabilities along with the availability, scalability, and security required by production workloads. The service offers built-in integrations with Kibana, Logstash, and AWS services including Amazon Kinesis Firehose, AWS Lambda, and Amazon CloudWatch so that you can go from raw data to actionable insights quickly. For more information on the elastic cache service, please refer to the below link:

- <https://aws.amazon.com/elasticsearch-service/>

#### NEW QUESTION 123

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

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



down the Dockerfile from S3 and run it when new instances launch.

D. Create a new Auto Scaling Launch Configuration with UserData scripts configured to pull the latest code at all times.

**Answer: B**

**Explanation:**

Since the time required to spin up an instance is required to be fast, its better to create an AMI rather than use User Data. When you use User Data, the script will be

run during boot up, and hence this will be slower.

An Amazon Machine Image (AMI) provides the information required to launch an instance, which is a virtual server in the cloud. You specify an AMI when you launch

an instance, and you can launch as many instances from the AMI as you need. You can also launch instances from as many different AMIs as you need.

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

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

**NEW QUESTION 127**

You need to scale an RDS deployment. You are operating at 10% writes and 90% reads, based on your logging. How best can you scale this in a simple way?

- A. Create a second master RDS instance and peer the RDS groups.
- B. Cache all the database responses on the read side with CloudFront.
- C. Create read replicas for RDS since the load is mostly reads.
- D. Create a Multi-AZ RDS installs and route read traffic to standby.

**Answer: C**

**Explanation:**

Amazon RDS Read Replicas provide enhanced performance and durability for database (DB) instances. This replication feature makes it easy to elastically scale out beyond the capacity constraints of a single DB Instance for read-heavy database workloads. You can create one or more replicas of a given source DB Instance and serve high-volume application read traffic from multiple copies of your data, thereby increasing aggregate read throughput. Read replicas can also be promoted when needed to become standalone DB instances.

Option A is invalid because you would need to maintain the synchronization yourself with a secondary instance.

Option B is invalid because you are introducing another layer unnecessarily when you already have read replica's Option D is invalid because you only use this for Standby's

For more information on Read Replica's, please refer to the below link: <https://aws.amazon.com/rds/details/read-replicas/>

**NEW QUESTION 132**

You are building a game high score table in DynamoDB. You will store each user's highest score for each game, with many games, all of which have relatively similar usage levels and numbers of players. You need to be able to look up the highest score for any game. What's the best DynamoDB key structure?

- A. HighestScore as the hash/only key.
- B. GameID as the hash key, HighestScore as the range key
- C. GameID as the hash/only key.
- D. GameID as the hash/only key.

**Answer: B**

**Explanation:**

It always best to choose the hash key as the column that will have a wide range of values. This is also given in the AWS documentation

Choosing a Partition Key

The following table compares some common partition key schemas for provisioned throughput efficiency:

Partition key value	Uniformity
User ID, where the application has many users.	Good
Status code, where there are only a few possible status codes.	Bad
Item creation date, rounded to the nearest time period (e.g. day, hour, minute)	Bad
Device ID, where each device accesses data at relatively similar intervals	Good
Device ID, where even if there are a lot of devices being tracked, one is by far more popular than all the others.	Bad

Next since you need to sort by the Highest Score, you need to use that as the sort key For more information on Table Guidelines, please visit the below URL:

- <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GuidelinesForTables.html>

**NEW QUESTION 135**

You need to run a very large batch data processingjob one time per day. The source data exists

entirely in S3, and the output of the processingjob should also be written to S3 when finished. If you need to version control this processingjob and all setup and teardown logic for the system, what approach should you use?.

- A. Model an AWSEMRjob in AWS Elastic Beanstalk.
- B. Model an AWSEMRjob in AWS CloudFormation.
- C. Model an AWS EMRjob in AWS OpsWorks.
- D. Model an AWS EMRjob in AWS CLI Composer.

**Answer: B**

**Explanation:**

With AWS Cloud Formation, you can update the properties for resources in your existing stacks.

These changes can range from simple configuration changes, such

as updating the alarm threshold on a Cloud Watch alarm, to more complex changes, such as updating the Amazon Machine Image (AMI) running on an Amazon EC2

instance. Many of the AWS resources in a template can be updated, and we continue to add support for more.

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

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

#### NEW QUESTION 136

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

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

**Answer:** B

#### Explanation:

The CBS Snapshots are stored in S3, so if you have a script which deploys EC2 Instances, the CBS volumes need to be constructed from snapshots stored in S3.

You can back up the data on your Amazon CBS volumes to Amazon S3 by taking point-in-time snapshots. Snapshots are incremental backups, which means that only the blocks on the device that have changed after your most recent snapshot are saved. This minimizes the time required to create the snapshot and saves on storage costs by not duplicating data. When you delete a snapshot, only the data unique to that snapshot is removed. Each snapshot contains all of the information needed to restore your data (from the moment when the snapshot was taken) to a new CBS volume. For more information on CBS Snapshots, please visit the below URL:

- <http://docs.aws.amazon.com/AWSC2/latest/UserGuide/CBSSnapshots.html>

#### NEW QUESTION 137

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-alloc-tags.html>

#### NEW QUESTION 138

You have an asynchronous processing application using an Auto Scaling Group and an SQS Queue. The Auto Scaling Group scales according to the depth of the job queue. The completion velocity of the jobs has gone down, the Auto Scaling Group size has maxed out, but the inbound job velocity did not increase. What is a possible issue?

- A. Some of the new jobs coming in are malformed and unprocessable.
- B. The routing tables changed and none of the workers can process events anymore.
- C. Someone changed the IAM Role Policy on the instances in the worker group and broke permissions to access the queue.
- D. The scaling metric is not functioning correctly.

**Answer:** A

#### Explanation:

This question is more on the grounds of validating each option

Option B is invalid, because the Route table would have an effect on all worker processes and no jobs would have been completed.

Option C is invalid because if the IAM Role was invalid then no jobs would be completed.

Option D is invalid because the scaling is happening, it's just that the jobs are not getting completed. For more information on Scaling on Demand, please visit the below URL:

- <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scale-based-on-demand.html>

#### NEW QUESTION 141

What is required to achieve gigabit network throughput on EC2? You already selected cluster- compute, 10GB instances with enhanced networking, and your workload is already network-bound, but you are not seeing 10 gigabit speeds.

- A. Enable bi-directional networking on your servers, so packets are non-blocking in both directions and there's no switching overhead.
- B. Ensure the instances are in different VPCs so you don't saturate the Internet Gateway on any one VPC.
- C. Select PIOPS for your drives and mount several, so you can provision sufficient disk throughput.
- D. Use a placement group for your instances so the instances are physically near each other in the same Availability Zone.

**Answer:** D

**Explanation:**

A placement group is a logical grouping of instances within a single Availability Zone. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. To provide the lowest latency, and the highest packet-per-second network performance for your placement group, choose an instance type that supports enhanced networking. For more information on Placement Groups, please visit the below URL:  
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

**NEW QUESTION 144**

You need to deploy a new application version to production. Because the deployment is high-risk, you need to roll the new version out to users over a number of hours, to make sure everything is working correctly. You need to be able to control the proportion of users seeing the new version of the application down to the percentage point. You use ELB and EC2 with Auto Scaling Groups and custom AMIs with your code pre-installed assigned to Launch Configurations. There are no data base- level changes during your deployment. You have been told you cannot spend too much money, so you must not increase the number of EC2 instances much at all during the deployment, but you also need to be able to switch back to the original version of code quickly if something goes wrong. What is the best way to meet these requirements?


- A. Create a second ELB, Auto Scaling Launch Configuration, and Auto Scaling Group using the Launch Configuratio
- B. Create AMIs with all code pre-installe
- C. Assign the new AMI to the second Auto Scaling Launch Configuratio
- D. Use Route53 Weighted Round Robin Records to adjust the proportion of traffic hitting the two ELBs.S
- E. Use the Blue-Green deployment method to enable the fastest possible rollback if neede
- F. Create a full second stack of instances and cut the DNS over to the new stack of instances, and change the DNS back if a rollback is needed.
- G. Create AMIs with all code pre-installe
- H. Assign the new AMI to the Auto Scaling Launch Configuration, to replace the old on
- I. Gradually terminate instances running the old code (launched with the old Launch Configuration) and allow the new AMIs to boot to adjust the traffic balance to the new cod
- J. On rollback, reverse the process by doing the same thing, but changing the AMI on the Launch Config back to the original code.
- K. Migrate to use AWS Elastic Beanstal
- L. Use the established and well-tested Rolling Deployment setting AWS provides on the new Application Environment, publishing a zip bundle of the new code and adjusting the wait period to spread the deployment over tim
- M. Re-deploy the old code bundle to rollback if needed.

**Answer:** A

**Explanation:**

This is an example of a Blue Green Deployment

You can shift traffic all at once or you can do a weighted distribution. With Amazon Route 53, you can define a percentage of traffic to go to the green environment and gradually update the weights until the green environment carries the full production traffic. A weighted distribution provides the ability to perform canary analysis where a small percentage of production traffic is introduced to a new environment. You can test the new code and monitor for errors, limiting the blast radius if any issues are encountered. It also allows the green environment to scale out to support the full production load if you're using Elastic Load Balancing



For more information on Blue Green Deployments, please visit the below URL:

- [https://dOawsstatic.com/whitepapers/AWS\\_Blue\\_Green\\_Deployments.pdf](https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf)

**NEW QUESTION 148**

You are building a mobile app for consumers to post cat pictures online. You will be storing the images in AWS S3. You want to run the system very cheaply and simply. Which one of these options allows you to build a photo sharing application with the right authentication/authorization implementation.

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

**Answer:** A

**Explanation:**

Amazon Cognito lets you easily add user sign-up and sign-in and manage permissions for your mobile and web apps. You can create your own user directory within Amazon Cognito. You can also choose to authenticate users through social identity providers such as Facebook, Twitter, or Amazon; with SAML identity



solutions; or by using your own identity system. In addition, Amazon Cognito enables you to save data locally on users' devices, allowing your applications to work even when the devices are offline. You can then synchronize data across users' devices so that their app experience remains consistent regardless of the device they use.

For more information on AWS Cognito, please visit the below URL:

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

#### NEW QUESTION 153

Your team wants to begin practicing continuous delivery using CloudFormation, to enable automated builds and deploys of whole, versioned stacks or stack layers. You have a 3-tier, mission-critical system. Which of the following is NOT a best practice for using CloudFormation in a continuous delivery environment?

- A. Use the AWS CloudFormation ValidateTemplate call before publishing changes to AWS.
- B. Model your stack in one template, so you can leverage CloudFormation's state management and dependency resolution to propagate all changes.
- C. Use CloudFormation to create brand new infrastructure for all stateless resources on each push, and run integration tests on that set of infrastructure.
- D. Parametrize the template and use Mappings to ensure your template works in multiple Regions.

**Answer: B**

#### Explanation:

Answer - B

Some of the best practices for Cloudformation are

- Created Nested stacks

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::Stackresource in your template to reference other templates.

- Reuse Templates

After you have your stacks and resources set up, you can reuse your templates to replicate your infrastructure in multiple environments. For example, you can create environments for development, testing, and production so that you can test changes before implementing them into production. To make templates reusable, use the parameters, mappings, and conditions sections so that you can customize your stacks when you create them. For example, for your development environments, you can specify a lower-cost instance type compared to your production environment, but all other configurations and settings remain the same. For more information on Cloudformation best practises, please visit the below URL:

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

#### NEW QUESTION 157

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

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

**Answer: D**

#### Explanation:

Failover routing lets you route traffic to a resource when the resource is healthy or to a different resource when the first resource is unhealthy. The primary and secondary resource record sets can route traffic to anything from an Amazon S3 bucket that is configured as a website to a complex tree of records.

For more information on Route53 Failover Routing, please visit the below URL:

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

#### NEW QUESTION 159

Your CTO thinks your AWS account was hacked. What is the only way to know for certain if there was unauthorized access and what they did, assuming your hackers are very sophisticated AWS engineers and doing everything they can to cover their tracks?

- A. Use CloudTrail Log File Integrity Validation.
- B. Use AWS Config SNS Subscriptions and process events in real time.
- C. Use CloudTrail backed up to AWS S3 and Glacier.
- D. Use AWS Config Timeline forensics.

**Answer: A**

#### Explanation:

To determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it, you can use CloudTrail log file integrity validation. This feature is built using industry standard algorithms: SHA-256 for hashing and SHA-256 with RSA for digital signing. This makes it computationally infeasible to modify, delete or forge CloudTrail log files without detection. You can use the AWS CLI to validate the files in the location where CloudTrail delivered them

Validated log files are invaluable in security and forensic investigations. For example, a validated log file enables you to assert positively that the log file itself has not changed, or that particular user credentials performed specific API activity. The CloudTrail log file integrity validation process also lets you know if a log file has been deleted or changed, or assert positively that no log files were delivered to your account during a given period of time.

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

[http://docs.aws.a mazon.com/awsccloudtrail/latest/userguide/cloudtrai l-log-file-validation- intro.html](http://docs.aws.amazon.com/awsccloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html)

#### NEW QUESTION 160

You need to grant a vendor access to your AWS account. They need to be able to read protected messages in a private S3 bucket at their leisure. They also use AWS. What is the best way to accomplish this?

- A. Create an IAM User with API Access Key

- B. Grant the User permissions to access the bucket
- C. Give the vendor the AWS Access Key ID and AWS Secret Access Key for the User.
- D. Create an EC2 Instance Profile on your account
- E. Grant the associated IAM role full access to the bucket
- F. Start an EC2 instance with this Profile and give SSH access to the instance to the vendor.
- G. Create a cross-account IAM Role with permission to access the bucket, and grant permission to use the Role to the vendor AWS account.
- D- Generate a signed S3 PUT URL and a signed S3 GET URL, both with wildcard values and 2 year duration
- H. Pass the URLs to the vendor.

**Answer:** C

**Explanation:**

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 Access, please visit the below URL:

- <https://aws.amazon.com/blogs/security/tag/cross-account-access/>

**NEW QUESTION 164**

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 165**

You have an application hosted in AWS, which sits on EC2 Instances behind an Elastic Load Balancer. You have added a new feature to your application and are now receiving complaints from users that the site has a slow response. Which of the below actions can you carry out to help you pinpoint the issue?

- A. Use CloudTrail to log all the API calls, and then traverse the log files to locate the issue
- B. Use CloudWatch, monitor the CPU utilization to see the times when the CPU peaked
- C. Review the Elastic Load Balancer logs
- D. Create some custom CloudWatch metrics which are pertinent to the key features of your application

**Answer:** D

**Explanation:**

Since the issue is occurring after the new feature has been added, it could be relevant to the new feature.

Enabling CloudTrail will just monitor all the API calls of all services and will not benefit the cause.

The monitoring of CPU utilization will just verify that there is an issue but will not help pinpoint the issue.

The Elastic Load Balancer logs will also just verify that there is an issue but will not help pinpoint the issue.

For more information on custom CloudWatch metrics, please refer to the below link:

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

**NEW QUESTION 168**

You currently have EC2 Instances hosting an application. These instances are part of an AutoScaling Group. You now want to change the instance type of the EC2 Instances. How can you manage the deployment with the least amount of downtime?

- A. Terminate the existing Auto Scaling group
- B. Create a new launch configuration with the new Instance type
- C. Attach that to the new AutoScaling Group.
- D. Use the AutoScalingRollingUpdate policy on CloudFormation Template Auto Scaling group
- E. Use the Rolling Update feature which is available for EC2 Instances.
- F. Manually terminate the instances, launch new instances with the new instance type and attach them to the AutoScaling group

**Answer:** B

**Explanation:**

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 CloudFormation 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.

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

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

**NEW QUESTION 173**

When creating an Elastic Beanstalk environment using the Wizard, what are the 3 configuration options presented to you

- A. Choosing the type of Environment- Web or Worker environment
- B. Choosing the platform type- Node.js, IIS, etc
- C. Choosing the type of Notification - SNS or SQS
- D. Choosing whether you want a highly available environment or not

**Answer:** ABD

**Explanation:**

The below screens are what are presented to you when creating an Elastic Beanstalk environment



The high availability preset includes a load balancer; the low cost preset does not. For more information on the configuration settings, please refer to the below link:  
<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environments-create-wizard.html>

**NEW QUESTION 174**

An EC2 instance has failed a health check. What will the ELB do?

- A. The ELB will terminate the instance
- B. The ELB stops sending traffic to the instance that failed its health check
- C. The ELB does nothing
- D. The ELB will replace the instance

**Answer:** B

**Explanation:**

The AWS Documentation mentions

The load balancer routes requests only to the healthy instances. When the load balancer determines that an instance is unhealthy, it stops routing requests to that instance. The load balancer resumes routing requests to the instance when it has been restored to a healthy state.

For more information on ELB health checks, please refer to the below link: <http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-healthchecks.html>

**NEW QUESTION 175**

Which of the following services can be used in conjunction with Cloudwatch Logs. Choose the 3 most viable services from the options given below

- A. Amazon Kinesis
- B. Amazon S3
- C. Amazon SQS
- D. Amazon Lambda

**Answer:** ABD

**Explanation:**

The AWS Documentation the following products which can be integrated with Cloudwatch logs

- 1) Amazon Kinesis - Here data can be fed for real time analysis
- 2) Amazon S3 - You can use CloudWatch Logs to store your log data in highly durable storage such as S3.
- 3) Amazon Lambda - Lambda functions can be designed to work with Cloudwatch log. For more information on Cloudwatch Logs, please refer to the below link:  
[link:http://docs^ws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html](http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html)

**NEW QUESTION 178**

You currently have an application with an Auto Scaling group with an Elastic Load Balancer configured in AWS. After deployment users are complaining of slow response time for your application. Which of the following can be used as a start to diagnose the issue

- A. Use Cloudwatch to monitor the HealthyHostCount metric
- B. Use Cloudwatch to monitor the ELB latency
- C. Use Cloudwatch to monitor the CPU Utilization
- D. Use Cloudwatch to monitor the Memory Utilization

**Answer:** B

**Explanation:**



High latency on the ELB side can be caused by several factors, such as:

- Network connectivity
- ELB configuration
- Backend web application server issues

For more information on ELB latency, please refer to the below link:

- <https://aws.amazon.com/premiumsupport/knowledge-center/elb-latency-troubleshooting/>

#### NEW QUESTION 179

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. RedHatOpensift 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](http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_ecs.html)

#### NEW QUESTION 181

You need to investigate one of the instances which is part of your Autoscaling Group. How would you implement this.

- A. Suspend the AZRebalance process so that Autoscaling will not terminate the instance
- B. Put the instance in a standby state
- C. Put the instance in a InService state
- D. Suspend the AddToLoadBalancer process

**Answer: B**

#### Explanation:

The AWS Documentation mentions

Auto Scaling enables you to put an instance that is in the InService state into the Standbystate, update or troubleshoot the instance, and then return the instance to service. Instances that are on standby are still part of the Auto Scaling group, but they do not actively handle application traffic.

For more information on the standby state please refer to the below link:

- <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-enter-exit-standby.html>

#### NEW QUESTION 184

Which of the following services can be used to implement DevOps in your company.

- A. AWS Elastic Beanstalk
- B. AWSOpswork
- C. AWS Cloudformation
- D. All of the above

**Answer: D**

#### Explanation:

All of the services can be used to implement Devops in your company

1) AWS Elastic Beanstalk, an easy-to-use service for deploying and scaling web applications and services developed with Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker on servers such as Apache, Nginx, Passenger, and IIS.

2) AWS Ops Works, a configuration management service that helps you configure and operate applications of all shapes and sizes using Chef

3) AWS Cloud Formation, which is an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable fashion.

For more information on AWS Devops please refer to the below link:

- <http://docs.aws.amazon.com/devops/latest/gsg/welcome.html>

#### NEW QUESTION 187

You need to deploy a multi-container Docker environment on to Elastic beanstalk. Which of the following files can be used to deploy a set of Docker containers to Elastic beanstalk

- A. Dockerfile
- B. DockerMultifile
- C. Dockerrun.aws.json
- D. Dockerrun

**Answer: C**

#### Explanation:

The AWS Documentation specifies

A Dockerrun.aws.json file is an Elastic Beanstalk-specific JSON file that describes how to deploy a set of Docker containers as an Elastic Beanstalk application.

You can use aDockerrun.aws.json file for a multicontainer Docker environment.

Dockerrun.aws.json describes the containers to deploy to each container instance in the environment as well as the data volumes to create on the host instance for the containers to mount.

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

[http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create\\_deploy\\_docker\\_v2config.html](http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_v2config.html)

#### NEW QUESTION 190

.....

## 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