

## DOP-C01 Dumps

### AWS Certified DevOps Engineer- Professional

<https://www.certleader.com/DOP-C01-dumps.html>



**NEW QUESTION 1**

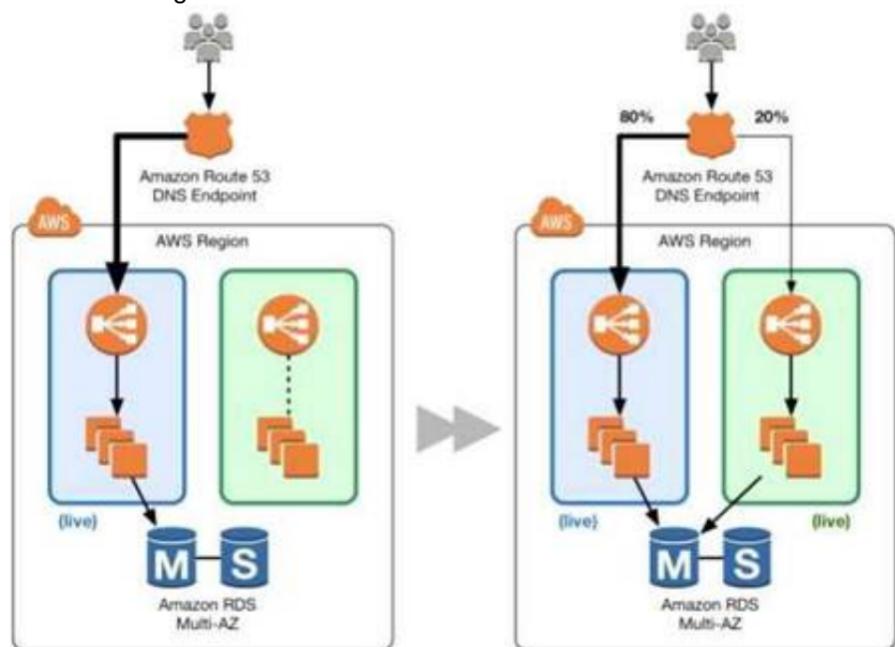
Your application is currently running on Amazon EC2 instances behind a load balancer. Your management has decided to use a Blue/Green deployment strategy. How should you implement this for each deployment?

- A. Set up Amazon Route 53 health checks to fail over from any Amazon EC2 instance that is currently being deployed to.
- B. Using AWS CloudFormation, create a test stack for validating the code, and then deploy the code to each production Amazon EC2 instance.
- C. Create a new load balancer with new Amazon EC2 instances, carry out the deployment, and then switch DNS over to the new load balancer using Amazon Route 53 after testing.
- D. Launch more Amazon EC2 instances to ensure high availability, de-register each Amazon EC2 instance from the load balancer, upgrade it, and test it, and then register it again with the load balancer.

**Answer: C**

**Explanation:**

The below diagram shows how this can be done



- 1) First create a new ELB which will be used to point to the new production changes.
  - 2) Use the Weighted Route policy for Route53 to distribute the traffic to the 2 ELB's based on a 80- 20% traffic scenario. This is the normal case, the % can be changed based on the requirement.
  - 3) Finally when all changes have been tested, Route53 can be set to 100% for the new ELB.
- Option A is incorrect because this is a failover scenario and cannot be used for Blue green deployments. In Blue Green deployments, you need to have 2 environments running side by side. Option B is incorrect, because you need to have a production stack with the changes which will run side by side. Option D is incorrect because this is not a blue green deployment scenario. You cannot control which users will go to the new EC2 instances. 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 2**

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 3**

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 4

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 5

Your application stores sensitive information on an EBS volume attached to your EC2 instance. How can you protect your information? Choose two answers from the options given below

- A. Unmount the EBS volume, take a snapshot and encrypt the snapshot
- B. Re-mount the Amazon EBS volume
- C. It is not possible to encrypt an EBS volume, you must use a lifecycle policy to transfer data to S3 for encryption.
- D. Copy the unencrypted snapshot and check the box to encrypt the new snapshot
- E. Volumes restored from this encrypted snapshot will also be encrypted.
- F. Create and mount a new, encrypted Amazon EBS volume
- G. Move the data to the new volume
- H. Delete the old Amazon EBS volume

**Answer: CD**

#### Explanation:

These steps are given in the AWS documentation

To migrate data between encrypted and unencrypted volumes

1) Create your destination volume (encrypted or unencrypted, depending on your need).

2) Attach the destination volume to the instance that hosts the data to migrate.

3) Make the destination volume available by following the procedures in Making an Amazon EBS Volume Available for Use. For Linux instances, you can create a mount point at /mnt/destination and mount the destination volume there.

4) Copy the data from your source directory to the destination volume. It may be most convenient to use a bulk-copy utility for this.

To encrypt a volume's data by means of snapshot copying

1) Create a snapshot of your unencrypted EBS volume. This snapshot is also unencrypted.

2) Copy the snapshot while applying encryption parameters. The resulting target snapshot is encrypted.

3) Restore the encrypted snapshot to a new volume, which is also encrypted.

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

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.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 sent to all the on-call teams of the company. With Option B & C, since we are not configuring the SNS topic for each on-call team the notifications will be sent to all the on-call teams. Hence these 2 options are invalid. For more information on setting up notifications, please refer to the below document link: from AWS [http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US\\_SetupSNS.html](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 using an AWS CloudFormation template
- 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 template
- D. Spin up a second AWS CloudFormation stack
- 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 Scaling
- G. Setup your Auto Scaling group policies to scale based on the number of requests per second as well as the current customer load time
- H. Re-deploy your application using an Auto Scaling template
- 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 EC2 instances available to handle the load for your application. You create collections of

EC2 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 your group 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 your group 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**

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 10**

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 10**

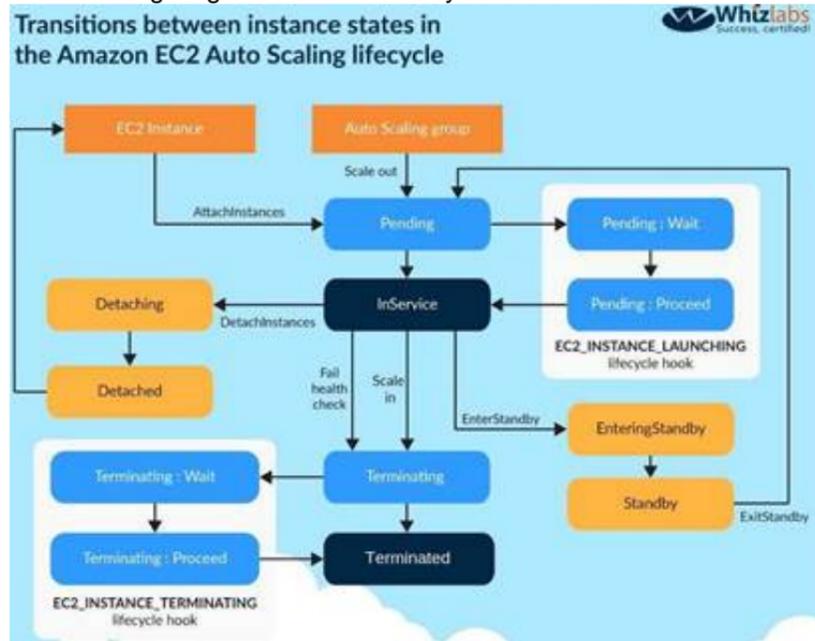
You have an Auto Scaling group of Instances that processes messages from an Amazon Simple Queue Service (SQS) queue. The group scales on the size of the queue. Processing involves calling a third-party web service. The web service is complaining about the number of failed and repeated calls it is receiving 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/AutoScalingGroup\\_lifecycle.htm](http://docs.aws.amazon.com/autoscaling/latest/userguide/AutoScalingGroup_lifecycle.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 11**

You are doing a load testing exercise on your application hosted on AWS. While testing your Amazon RDS MySQL DB instance, you notice that when you hit 100% CPU utilization on it, your application becomes non-responsive. Your application is read-heavy. What are methods to scale your data tier to meet the application's needs? Choose three answers from the options given below

- A. Add Amazon RDS DB read replicas, and have your application direct read queries to them.
- B. Add your Amazon RDS DB instance to an Auto Scaling group and configure your Cloud Watch metric based on CPU utilization.

- C. Use an Amazon SQS queue to throttle data going to the Amazon RDS DB instance.
- D. Use ElastiCache in front of your Amazon RDS DB to cache common queries.
- E. Shard your data set among multiple Amazon RDS DB instances.
- F. Enable Multi-AZ for your Amazon RDS DB instance.

**Answer:** ADE

**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

For more information on Read Replica's please refer to the below link:

? <https://aws.amazon.com/rds/details/read-replicas/>

Sharding is a common concept to split data across multiple tables in a database For more information on sharding please refer to the below link:

<https://forums.aws.amazon.com/thread.jspa?messageID=203052>

Amazon ElastiCache is a web service that makes it easy to deploy, operate, and scale an in-memory data store or cache in the cloud. The service improves the performance of web applications by allowing you to retrieve information from fast, managed, in-memory data stores, instead of relying entirely on slower disk-based databases

Amazon ElastiCache is an in-memory key/value store that sits between your application and your database. Whenever your application requests data, it first makes the request to the ElastiCache cache. If the data exists in the cache and is current, ElastiCache returns the data to your application. If the data does not exist in the cache, or the data in the cache has expired, your application requests the data from your database which returns the data to your application. Your application then writes the data received from the database to the cache so it can be more quickly retrieved next time it is requested. For more information on ElastiCache please refer to the below link:

<https://aws.amazon.com/elasticache/>

Option B is not an ideal way to scale a database

Option C is not ideal to store the data which would go into a database because of the message size Option F is invalid because Multi-AZ feature is only a failover option

**NEW QUESTION 14**

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

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

**Answer:** A

**Explanation:**

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

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

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

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

**NEW QUESTION 19**

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 it 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 20**

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 EC2 creates snapshots of your instance's root volume and any other EBS volumes attached to your instance. New EBS 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 be 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 EBS initialization please refer to the below link:

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

**NEW QUESTION 24**

You have enabled Elastic Load Balancing HTTP health checking. After looking at the AWS Management Console, you see that all instances are passing health checks, but your customers are reporting that your site is not responding. What is the cause?

- A. The HTTP health checking system is misreporting due to latency in inter-instance metadata synchronization.
- B. The health check in place is not sufficiently evaluating the application function.
- C. The application is returning a positive health check too quickly for the AWS Management Console to respond.
- D. Latency in DNS resolution is interfering with Amazon EC2 metadata retrieval.

**Answer: B**

**Explanation:**

You need to have a custom health check which will evaluate the application functionality. It's not enough using the normal health checks. If the application functionality does not work and if you don't have custom health checks, the instances will still be deemed as healthy. 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 Auto Scaling health checks, please refer to the below document link: from AWS <http://docs.aws.amazon.com/autoscaling/latest/userguide/healthcheck.html>

**NEW QUESTION 26**

When an Auto Scaling group is running in Amazon Elastic Compute Cloud (EC2), your application rapidly scales up and down in response to load within a 10-minute window; however, after the load peaks, you begin to see problems in your configuration management system where previously terminated Amazon EC2 resources are still showing as active. What would be a reliable and efficient way to handle the cleanup of Amazon EC2 resources within your configuration management system? Choose two answers from the options given below.

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

**Answer: AD**

**Explanation:**

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

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

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

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

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

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

**NEW QUESTION 29**

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

- A. Create a Chef Recipe to update this configuration file, configure your AWS OpsWorks stack to use custom cookbooks, and assign this recipe to the Configure Lifecycle Event of the specific layer.
- B. Update this configuration file by writing a script to poll the AWS OpsWorks service API for new instances.
- C. Configure your base AMI to execute this script on Operating System startup.

- D. Create a Chef Recipe to update this configuration file, configure your AWS OpsWorks stack to use custom cookbooks, and assign this recipe to execute when instances are launched.
- E. Configure your AWS OpsWorks layer to use the AWS-provided recipe for distributed host configuration, and configure the instance hostname and file path parameters in your recipes settings.

**Answer:** A

**Explanation:**

Please check the following AWS DOCs which provides details on the scenario. Check the example of "configure".  
? <https://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-events.html> You can use the Configure Lifecycle event  
This event occurs on all of the stack's instances when one of the following occurs:

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

2. Toggle **Use custom Chef cookbooks** to **Yes**.

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

- [http://docs.aws.amazon.com/opsworks/latest/userguide/welcome\\_classic.html](http://docs.aws.amazon.com/opsworks/latest/userguide/welcome_classic.html)

**NEW QUESTION 30**

You have a web application that's developed in Node.js The code is hosted in Git repository. You want to now deploy this application to AWS. Which of the below 2 options can fulfil this requirement.

- A. Create an Elastic Beanstalk applicatio
- B. Create a Docker file to install Node.j
- C. Get the code from Gi
- D. Use the command "aws git.push" to deploy the application
- E. Create an AWS CloudFormation template which creates an instance with the AWS::EC2::Container resources typ
- F. With UserData, install Git to download the Node.js application and then set it up.
- G. Create a Docker file to install Node.j
- H. and gets the code from Gi
- I. Use the Dockerfile to perform the deployment on a new AWS Elastic Beanstalk applicatio
- J. S
- K. Create an AWS CloudFormation template which creates an instance with the AWS::EC2::Instance resource type and an AMI with Docker pre-installe
- L. With UserData, install Git to download the Node.js application and then set it up.

**Answer:** CD

**Explanation:**

Option A is invalid because there is no "awsgitpush" command  
Option B is invalid because there is no AWS::CC2::Container resource type.  
Clastic Beanstalk supports the deployment of web applications from Docker containers. With Docker containers, you can define your own runtime environment. You can choose your own platform, programming language, and any application dependencies (such as package managers or tools), that aren't supported by other platforms. Docker containers are self-contained and include all the configuration information and software your web application requires to run.

For more information on Docker and Clastic beanstalk please refer to the below link:

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

When you launch an instance in Amazon CC2, you have the option of passing user data to the instance that can be used to perform common automated configuration tasks and even run scripts after the instance starts. You can pass two types of user data to Amazon CC2: shell scripts and cloud- init directives. You can also pass this data into the launch wizard as plain text, as a file (this is useful for launching instances using the command line tools), or as base64-encoded text (for API calls). For more information on Cc2 User data please refer to the below link:

- <http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/user-data.htm>

Note: "git aws.push" with CB CLI 2.x - see a forum thread at <https://forums.aws.amazon.com/thread.jspa?messageID=583202#jive-message-582979>. Basically, this is a predecessor to the newer "eb deploy" command in CB CLI 31. This question kept in order to be consistent with exam.

**NEW QUESTION 32**

Your company develops a variety of web applications using many platforms and programming languages with different application dependencies. Each application must be developed and deployed quickly and be highly available to satisfy your business requirements. Which of the following methods should you use to deploy these applications rapidly?

- A. Develop the applications in Docker containers, and then deploy them to Elastic Beanstalk environments with Auto Scaling and Elastic Load Balancing.

- B. Use the AWS CloudFormation Docker import service to build and deploy the applications with high availability in multiple Availability Zones.
- C. Develop each application's code in DynamoDB, and then use hooks to deploy it to Elastic Beanstalk environments with Auto Scaling and Elastic Load Balancing.
- D. Store each application's code in a Git repository, develop custom package repository managers for each application's dependencies, and deploy to AWS OpsWorks in multiple Availability Zones.

**Answer:** A

**Explanation:**

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

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

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

**NEW QUESTION 34**

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

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

**Answer:** A

**Explanation:**

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

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

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

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

**NEW QUESTION 37**

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.

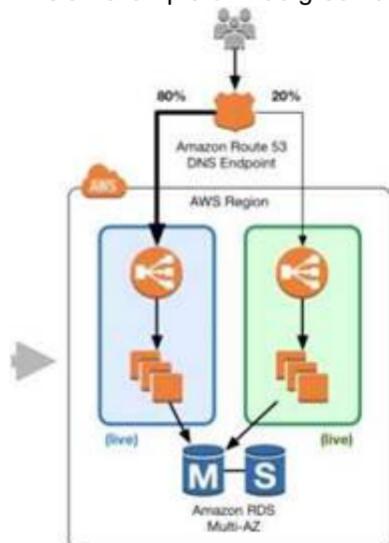
Which combination of services would give you the ability to control traffic between different deployed versions of your application?

- A. Create one AWS Elastic Beanstalk application and all AWS resources (using configuration files inside the application source bundle) for each web application
- B. New versions would be deployed using Elastic Beanstalk environments and using the Swap URLs feature.
- C. Using AWS CloudFormation templates, create one Elastic Beanstalk application and all AWS resources (in the same template) for each web application
- D. 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. >/
- E. Using AWS CloudFormation templates, create one Elastic Beanstalk application and all AWS resources (in the same template) for each web application
- F. 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.
- G. Create one Elastic Beanstalk application and all AWS resources (using configuration files inside the application source bundle) for each web application
- H. 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 40**

You've been tasked with improving the current deployment process by making it easier to deploy and reducing the time it takes. You have been tasked with creating a continuous integration (CI) pipeline that can build AMI'S. Which of the below is the best manner to get this done. Assume that at max your development team will be deploying builds 5 times a week.

- A. Use a dedicated EC2 instance with an EBS Volum
- B. Download and configure the code and then crate an AMI out of that.
- C. Use OpsWorks to launch an EBS-backed instance, then use a recipe to bootstrap the instance, and then have the CI system use the CreateImage API call to make an AMI from it.
- D. Upload the code and dependencies to Amazon S3, launch an instance, download the package from Amazon S3, then create the AMI with the CreateSnapshot API call
- E. Have the CI system launch a new instance, then bootstrap the code and dependencies on that instance, and create an AMI using the CreateImage API call.

**Answer:** D

**Explanation:**

Since the number of calls is just a few times a week, there are many open source systems such as Jenkins which can be used as CI based systems. Jenkins can be used as an extensible automation server, Jenkins can be used as a simple CI server or turned into the continuous delivery hub for any project. For more information on the Jenkins CI tool please refer to the below link:

- <https://jenkins.io/>

Option A and C are partially correct, but since you just have 5 deployments per week, having separate instances which consume costs is not required. Option B is partially correct, but again having a separate system such as Opswork for such a low number of deployments is not required.

**NEW QUESTION 41**

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?

- A. Use CloudFormation templates
- B. Configuration Management Templates
- C. Saved Configurations
- D. 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 46**

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

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

**Answer:** AB

**Explanation:**

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

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

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

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

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

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

**NEW QUESTION 48**

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

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

**Answer:** C

**Explanation:**

With web identity federation, you don't need to create custom sign-in code or manage your own user identities. Instead, users of your app can sign in using a well-known identity provider (IdP) — such as Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible IdP, receive an authentication token, and then exchange that token for temporary security credentials in AWS that map to an IAM role with permissions to use the resources in your AWS

account. Using an IdP helps you keep your AWS account secure, because you don't have to embed and distribute long-term security credentials with your application. For more information on Web Identity Federation, please refer to the below document link: from AWS  
[http://docs.wsamazon.com/IAM/latest/UserGuide/id\\_roles\\_providers\\_oidc.html](http://docs.wsamazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html)

#### NEW QUESTION 51

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 1AM 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 1AM role to your Amazon EC2 instance, and use this 1AM role to access the Amazon RDS DB from your Amazon EC2 instances.

**Answer: D**

#### Explanation:

Creating and Using an 1AM Policy for 1AM Database Access

To allow an 1AM user or role to connect to your DB instance or DB cluster, you must create an 1AM policy. After that you attach the policy to an 1AM user or role.

Note

To learn more about 1AM policies, see Authentication and Access Control for Amazon RDS.

The following example policy allows an 1AM user to connect to a DB instance using 1AM database authentication.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "rds-db:connect"
      ],
      "Resource": [
        "arn:aws:rds-db:us-west-2:123456789012:dbuser:db-12ABC34DEFG5HIJ6KLMNOP78QR/jane_doe"
      ]
    }
  ]
}
```

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 1AM database authentication. They aren't valid in any other context.

1AM 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 1AM. 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 1AM 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 55

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

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

**Answer: B**

#### Explanation:

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

The AWS::AutoScaling::AutoScalingGroup resource supports an UpdatePolicy attribute. This is used to define how an Auto Scaling group resource is updated

when

an update to the Cloud Formation stack occurs. A common approach to updating an Auto Scaling group is to perform a rolling update, which is done by specifying the AutoScalingRollingUpdate policy. This retains the same Auto Scaling group and replaces old instances with new ones, according to the parameters specified.

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

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

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

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

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

#### NEW QUESTION 59

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 63

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 65

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 69

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 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 73**

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

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

**Answer: D**

**Explanation:**

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

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

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

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

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

**NEW QUESTION 76**

You are hired as the new head of operations for a SaaS company. Your CTO has asked you to make debugging any part of your entire operation simpler and as fast as possible. She complains that she has no idea what is going on in the complex, service-oriented architecture, because the developers just log to disk, and it's very hard to find errors in logs on so many services. How can you best meet this requirement and satisfy your CTO?

- A. Copy all log files into AWS S3 using a cron job on each instanc
- B. Use an S3 Notification Configuration on the PutBucket event and publish events to AWS Lambd
- C. Use the Lambda to analyze logs as soon as they come in and flag issues.
- D. Begin using CloudWatch Logs on every servic
- E. Stream all Log Groups into S3 object
- F. Use AWS EMR clusterjobs to perform adhoc MapReduce analysis and write new queries when needed.
- G. Copy all log files into AWS S3 using a cron job on each instanc
- H. Use an S3 Notification Configuration on the PutBucket event and publish events to AWS Kinesi
- I. Use Apache Spark on AWS EMR to perform at-scale stream processing queries on the log chunks and flag issues.
- J. Begin using CloudWatch Logs on every servic
- K. Stream all Log Groups into an AWS Elastic search Service Domain running Kibana 4 and perform log analysis on a search cluster.

**Answer: D**

**Explanation:**

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

Oasticsearch Service is a fully managed service that delivers Dasticsearch'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 Cloud Watch so that you can go from raw data to actionable insights quickly. For more information on Elastic Search, please refer to the below link:

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

**NEW QUESTION 77**

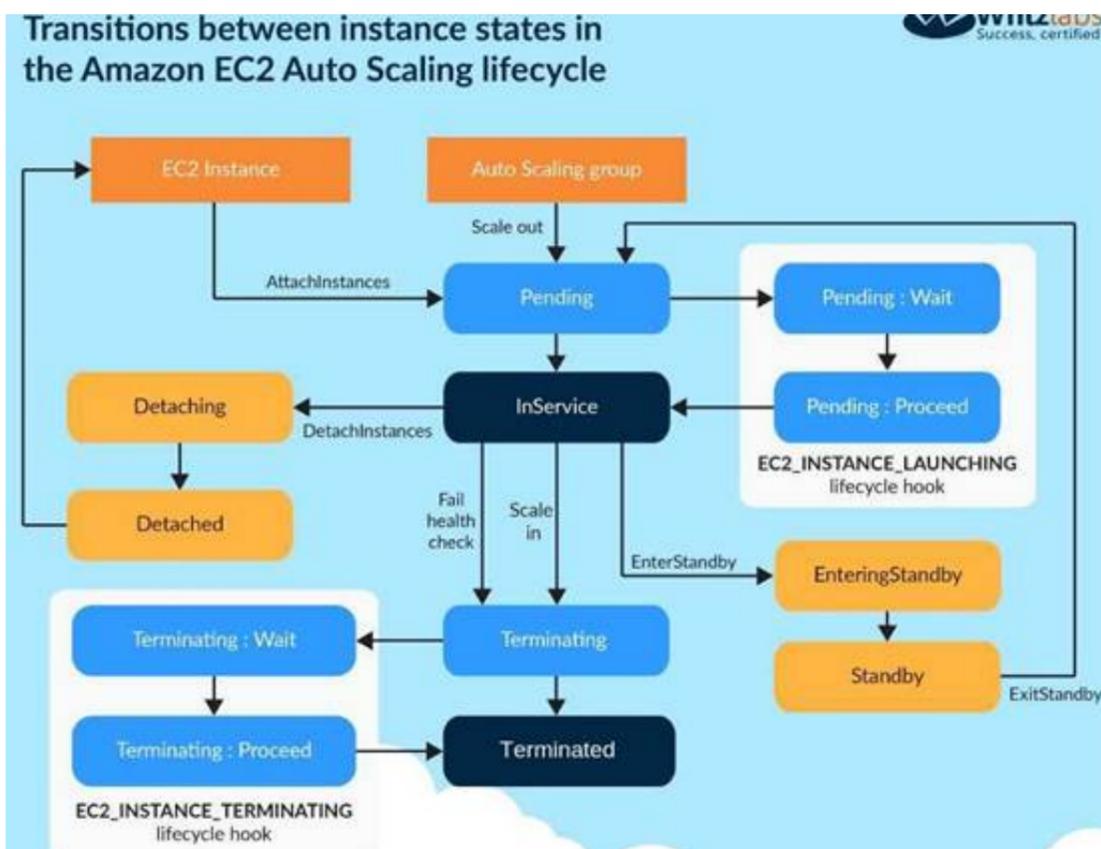
For AWS Auto Scaling, what is the first transition state an instance enters after leaving steady state when scaling in due to health check failure or decreased load?

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

**Answer: A**

**Explanation:**

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



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

**NEW QUESTION 80**

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.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 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 83**

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 84**

You have an application running on an Amazon EC2 instance and you are using IAM roles to securely access AWS Service APIs. How can you configure your application running on that instance to retrieve the API keys for use with the AWS SDKs?

- A. When assigning an EC2 IAM role to your instance in the console, in the "Chosen SDK" drop-down list, select the SDK that you are using, and the instance will configure the correct SDK on launch with the API keys.
- B. Within your application code, make a GET request to the IAM Service API to retrieve credentials for your user.
- C. When using AWS SDKs and Amazon EC2 roles, you do not have to explicitly retrieve API keys, because the SDK handles retrieving them from the Amazon EC2 Metadata service.
- D. Within your application code, configure the AWS SDK to get the API keys from environment variables, because assigning an Amazon EC2 role stores keys in environment variables on launch.

**Answer: C**

**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, 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>

**NEW QUESTION 85**

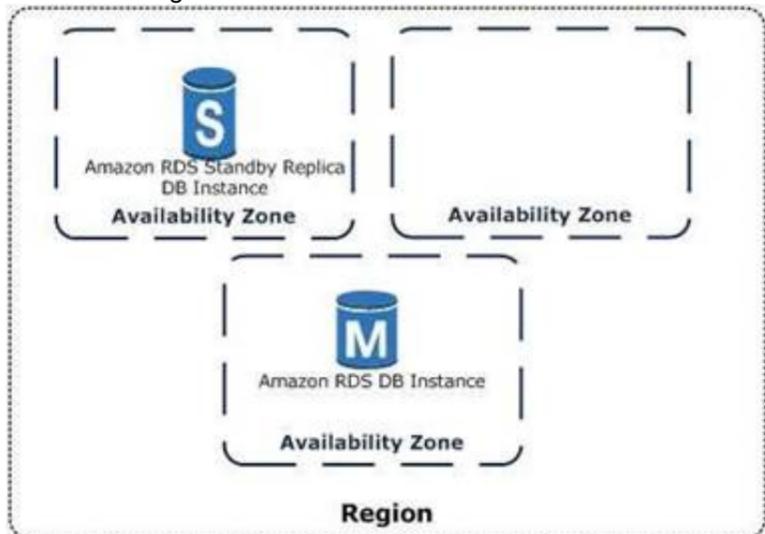
You are planning on using the Amazon RDS facility for Fault tolerance for your application. How does Amazon RDS Multi Availability Zone model work

- 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 90**

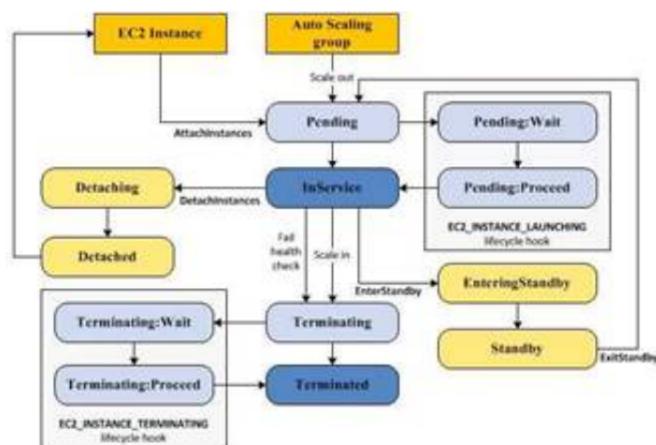
For AWS Auto Scaling, what is the first transition state an existing instance enters after leaving Standby state?

- A. Detaching
- B. Terminating:Wait
- C. Pending
- D. EnteringStandby

**Answer: C**

**Explanation:**

The below diagram shows the Lifecycle policy. When the stand-by state is exited, the next state is pending.



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

**NEW QUESTION 94**

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. AWSRedShift
- C. AWSEMR
- D. AWS DynamoDB

**Answer: A**

**Explanation:**

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

Oasticsearch Service is a fully managed service that delivers Dasticsearch'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 98**

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 101**

You are creating an application which stores extremely sensitive financial information. All information in the system must be encrypted at rest and in transit. Which of these is a violation of this policy?

- A. ELB SSL termination.
- B. ELB Using Proxy Protocol v1.
- C. CloudFront Viewer Protocol Policy set to HTTPS redirection.
- D. Telling S3 to use AES256 on the server-side.

**Answer: A**

**Explanation:**

If you use SSL termination, your servers will always get non-secure connections and will never know whether users used a more secure channel or not. If you are using Elastic beanstalk to configure the ELB, you can use the below article to ensure end to end encryption.

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/configuring-https-endtoend.html>

**NEW QUESTION 103**

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

- A. Use 9001 MTU instead of 1500 for Jumbo Frames, to raise packet body to packet overhead ratios.
- B. Segregate the instances into different peered VPCs while keeping them all in a placement group, so each one has its own Internet Gateway.

- C. Bake an AMI for the instances and relaunch, so the instances are fresh in the placement group and do not have noisy neighbors.  
D. Turn off SYN/ACK on your TCP stack or begin using UDP for higher throughput.

**Answer:** A

**Explanation:**

Jumbo frames allow more than 1500 bytes of data by increasing the payload size per packet, and thus increasing the percentage of the packet that is not packet overhead. Fewer packets are needed to send the same amount of usable data. However, outside of a given AWS region (CC2-Classic), a single VPC, or a VPC peering connection, you will experience a maximum path of 1500 MTU. VPN connections and traffic sent over an Internet gateway are limited to 1500 MTU. If packets are over 1500 bytes, they are fragmented, or they are dropped if the Don't Fragment flag is set in the IP header. For more information on Jumbo Frames, please visit the below URL:  
[http://docs.aws.amazon.com/AWSSCC2/latest/UserGuide/network\\_mtu.htm#jumbo\\_frame\\_instances](http://docs.aws.amazon.com/AWSSCC2/latest/UserGuide/network_mtu.htm#jumbo_frame_instances)

**NEW QUESTION 106**

You need to run a very large batch data processing job one time per day. The source data exists entirely in S3, and the output of the processing job should also be written to S3 when finished. If you need to version control this processing job 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 CloudFormation, 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 CloudWatch 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\\_stacks.html](http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/updates_stacks.html)

**NEW QUESTION 110**

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/AWSSCC2/latest/UserGuide/CBSnapshots.html>

**NEW QUESTION 115**

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 120**

You need your API backed by DynamoDB to stay online during a total regional AWS failure. You can tolerate a couple minutes of lag or slowness during a large failure event, but the system should recover with normal operation after those few minutes. What is a good approach?

- A. Set up DynamoDB cross-region replication in a master-standby configuration, with a single standby in another region
- B. Create an Auto Scaling Group behind an ELB in each of the two regions for your application layer in which DynamoDB is running in
- C. Add a Route53 Latency DNS Record with DNS Failover, using the ELBs in the two regions as the resource records.
- D. Set up a DynamoDB Global table
- E. Create an Auto Scaling Group behind an ELB in each of the two regions for your application layer in which the DynamoDB is running in
- F. Add a Route53 Latency DNS Record with DNS Failover, using the ELBs in the two regions as the resource records.
- G. Set up a DynamoDB Multi-Region table
- H. Create a cross-region ELB pointing to a cross-region Auto Scaling Group, and direct a Route53 Latency DNS Record with DNS Failover to the cross-region ELB.
- I. Set up DynamoDB cross-region replication in a master-standby configuration, with a single standby in another region
- J. Create a cross-region ELB pointing to a cross-region Auto Scaling Group, and direct a Route53 Latency DNS Record with DNS Failover to the cross-region ELB.

**Answer: B**

**Explanation:**

Updated based on latest AWS updates

Option A is invalid because using Latency based routing will send traffic on the region with the standby instance. This is an active/passive replication and you can't write to the standby table unless there is a failover. Answer A can work: only if you use a failover routing policy.

Option D is invalid because there is no concept of a cross-region CLB.

Amazon DynamoDB global tables provide a fully managed solution for deploying a multi-region, multi-master database, without having to build and maintain your own replication solution. When you create a global table, you specify the AWS regions where you want the table to be available. DynamoDB performs all of the necessary tasks to create identical tables in these regions, and propagate ongoing data changes to all of them.

For more information on DynamoDB Global Tables, please visit the below URL:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GlobalTables.html>

**NEW QUESTION 125**

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

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

**Answer: C**

**Explanation:**

Answer - C

You can monitor the progress of a stack update by viewing the stack's events. The console's Cvents tab displays each major step in the creation and update of the stack sorted by the time of each event with latest events on top. The start of the stack update process is marked with an UPDATE\_IN\_PROGRESS event for the stack. For more information on Monitoring your stack, please visit the below URL:

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

**NEW QUESTION 129**

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 biphase 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 134**

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

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

**Answer: A**

**Explanation:**

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

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

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

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

**NEW QUESTION 135**

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 policies
- 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-side
- E. Construct a custom build of the SDK and include S3 access in it.
- F. Create an AWS OAuth Service Domain and grant public signup and access to the domain
- 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 138**

You run a 2000-engineer organization. You are about to begin using AWS at a large scale for the first time. You want to integrate with your existing identity management system running on Microsoft Active Directory, because your organization is a power-user of Active Directory. How should you manage your AWS identities in the most simple manner?

- A. Use AWS Directory Service Simple AD.
- B. Use AWS Directory Service AD Connector.
- C. Use an Sync Domain running on AWS Directory Service.
- D. Use an AWS Directory Sync Domain running on AWS Lambda.

**Answer:** B

**Explanation:**

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

can support larger organizations of up to 5,000 users. Once set up, AD Connector offers the following benefits:

- Your end users and IT administrators can use their existing corporate credentials to log on to AWS applications such as Amazon Workspaces, Amazon WorkDocs, or Amazon WorkMail.
- You can manage AWS resources like Amazon EC2 instances or Amazon S3 buckets through IAM role-based access to the AWS Management Console.
- You can consistently enforce existing security policies (such as password expiration, password history, and account lockouts) whether users or IT administrators are accessing resources in your on-premises infrastructure or in the AWS Cloud.
- You can use AD Connector to enable multi-factor authentication by integrating with your existing RADIUS-based MFA infrastructure to provide an additional layer of security when users access AWS applications.

For more information on the AD Connector, please visit the below URL:

- [http://docs.aws.amazon.com/directoryservice/latest/admin-guide/directory\\_ad\\_connector.html](http://docs.aws.amazon.com/directoryservice/latest/admin-guide/directory_ad_connector.html)

**NEW QUESTION 143**

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

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

**Answer:** B

**Explanation:**

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

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

**NEW QUESTION 144**

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 149**

You currently have an application deployed via Elastic Beanstalk. You are now deploying a new application and have ensured that Elastic Beanstalk has detached the current instances and deployed and reattached new instances. But the new instances are still not receiving any sort of traffic. Why is this the case.

- A. The instances are of the wrong AMI, hence they are not being detected by the ELB.
- B. It takes time for the ELB to register the instances, hence there is a small timeframe before your instances can start receiving traffic
- C. You need to create a new Elastic Beanstalk application, because you cannot detach and then reattach instances to an ELB within an Elastic Beanstalk application
- D. The instances needed to be reattached before the new application version was deployed

**Answer: B**

**Explanation:**

Before the EC2 Instances can start receiving traffic, they will be checked via the health checks of the CLB. Once the health checks are successful, the EC2 Instance

will change its state to InService and then the EC2 Instances can start receiving traffic. 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 153**

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.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html](http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html)

**NEW QUESTION 156**

Which of the following features of the Elastic Beanstalk service will allow you to perform a Blue Green Deployment

- A. Rebuild Environment
- B. Swap Environment
- C. Swap URL's
- D. Environment Configuration

**Answer: C**

**Explanation:**

With the Swap url feature, you can keep a version of your environment ready. And when you are ready to cut over, you can just use the swap url feature to switch over

to your new environment

For more information on swap url feature, please refer to the below link:

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

**NEW QUESTION 157**

You have carried out a deployment using Elastic Beanstalk with All at once method, but the application is unavailable. What could be the reason for this

- A. You need to configure ELB along with Elastic Beanstalk
- B. You need to configure Route53 along with Elastic Beanstalk

- C. There will always be a few seconds of downtime before the application is available
- D. The cooldown period is not properly configured for Elastic Beanstalk

**Answer: C**

**Explanation:**

The AWS Documentation mentions

Because Elastic Beanstalk uses a drop-in upgrade process, there might be a few seconds of downtime. Use rolling deployments to minimize the effect of deployments on your production environments.

For more information on troubleshooting Elastic Beanstalk, please refer to the below link:

- <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/troubleshooting-deployments.html>
- <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-deploy-existing-version.html>

**NEW QUESTION 159**

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 163**

Which of the following is the right sequence of initial steps in the deployment of application revisions using Code Deploy

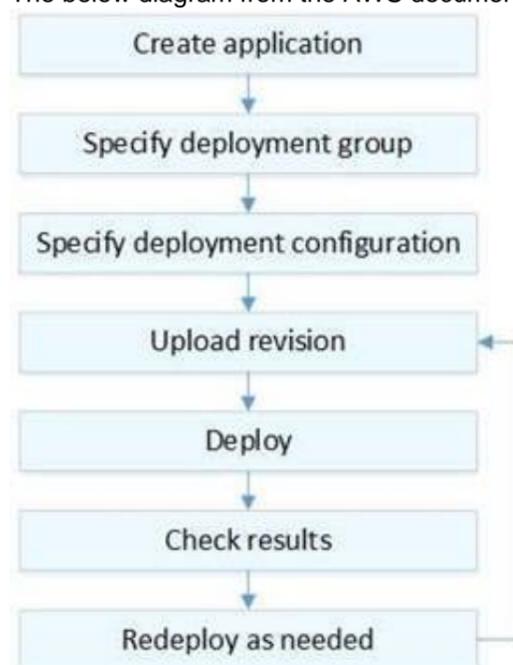
- 1) Specify deployment configuration
- 2) Upload revision
- 3) Create application
- 4) Specify deployment group

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

**Answer: C**

**Explanation:**

The below diagram from the AWS documentation shows the deployment steps



For more information on the deployment steps please refer to the below link:

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

**NEW QUESTION 166**

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 167**

Which of the following service can be used to provision ECS Cluster containing following components in an automated way:

- 1) Application Load Balancer for distributing traffic among various task instances running in EC2 Instances
- 2) Single task instance on each EC2 running as part of auto scaling group
- 3) Ability to support various types of deployment strategies

- A. SAM
- B. Opswork
- C. Elastic beanstalk
- D. CodeCommit

**Answer: C**

**Explanation:**

You can create docker environments that support multiple containers per Amazon EC2 instance with multi-container Docker platform for Elastic Beanstalk-Elastic Beanstalk uses Amazon Elastic Container Service (Amazon ECS) to coordinate container deployments to multi-container Docker environments. Amazon ECS provides tools to manage a cluster of instances running Docker containers. Elastic Beanstalk takes care of Amazon ECS tasks including cluster creation, task definition, and execution Please refer to the below AWS documentation: [https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create\\_deploy\\_docker\\_ecs.html](https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_ecs.html)

**NEW QUESTION 168**

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 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.

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 173**

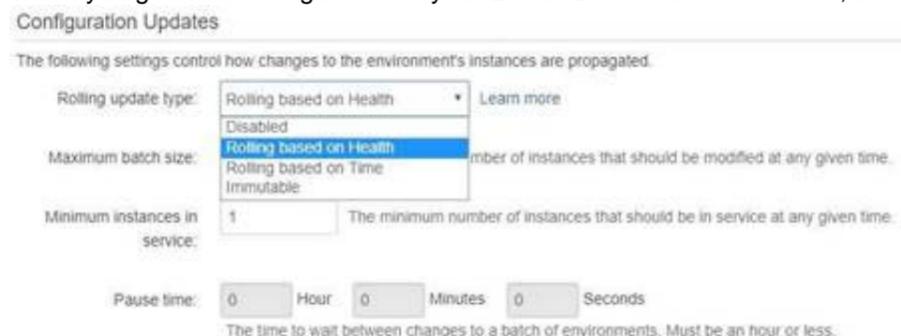
Which of the following is not a rolling type update which is present for Configuration Updates when it comes to the Elastic Beanstalk service

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

**Answer: B**

**Explanation:**

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



The AWS Documentation mentions

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

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

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

**NEW QUESTION 177**

You have the requirement to get a snapshot of the current configuration of the resources in your AWS Account. Which of the following services can be used for this purpose

- A. AWS CodeDeploy
- B. AWS Trusted Advisor
- C. AWSConfig
- D. AWSIAM

**Answer: C**

**Explanation:**

The AWS Documentation mentions the following With AWS Config, you can do the following:

- Evaluate your AWS resource configurations for desired settings.
- Get a snapshot of the current configurations of the supported resources that are associated with your AWS account.
- Retrieve configurations of one or more resources that exist in your account.
- Retrieve historical configurations of one or more resources.
- Receive a notification whenever a resource is created, modified, or deleted.
- View relationships between resources. For example, you might want to find all resources that use a particular security group. For more information on AWS Config, please visit the below URL: <http://docs.aws.amazon.com/config/latest/developerguide/WhatIsConfig.html>

**NEW QUESTION 182**

You have a video processing application hosted in AWS. The video's are uploaded by users onto the site. You have a program that is custom built to process those videos. The program is able to recover in case there are any failures when processing the videos. Which of the following mechanisms can be used to deploy the instances for carrying out the video processing activities, ensuring that the cost is kept at a minimum.

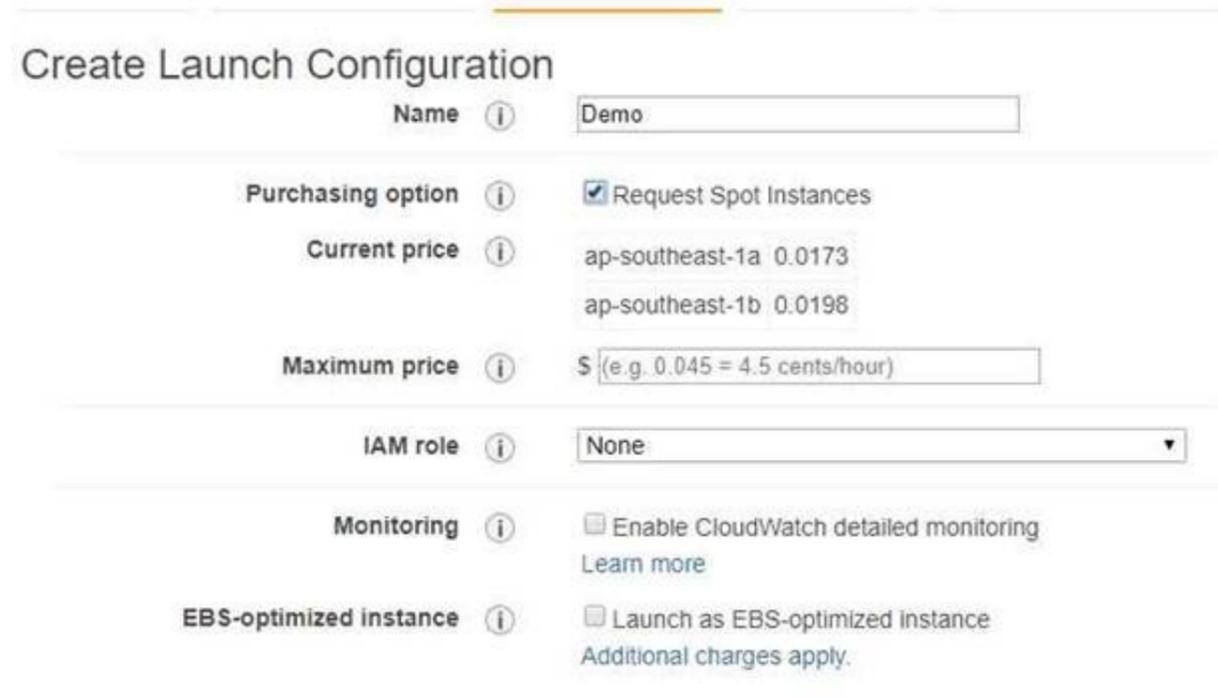
- A. Create a launch configuration with Reserved Instance
- B. Ensure the User Data section details the installation of the custom software
- C. Create an Autoscalinggroup with the launch configuration.
- D. Create a launch configuration with Spot Instance
- E. Ensure the User Data section details the installation of the custom software
- F. Create an Autoscalinggroupwith the launch configuration.
- G. Create a launch configuration with Dedicated Instance
- H. Ensure the User Data section details the installation of the custom software
- I. Create an Autoscaling group with the launch configuration.
- J. Create a launch configuration with On-Demand Instance
- K. Ensure the User Data section details the installation of the custom software
- L. Create an Autoscaling group with the launch configuration.

**Answer: B**

**Explanation:**

Since the application can recover from failures and cost is the priority, then Spot instances are the best bet for this requirement. The launch configuration has the facility to request for Spot Instances.

The below snapshot from the Launch configuration section shows that Spot Instances can be used for AutoScaling Groups.



► **Advanced Details**

For more information on Spot Instances and Autoscaling, please visit the below URL:

• <http://docs^ws.amazon.com/autoscaling/latest/userguide/US-SpotInstances.html>

**NEW QUESTION 184**

You currently have an Autoscalinggroup that has the following settings Min capacity-2

Desired capacity - 2 Maximum capacity - 2

Your launch configuration has AMI'S which are based on the t2.micro instance type. The application running on these instances are now experiencing issues and you have identified that the solution is to change the instance type of the instances running in the Autoscaling Group.

Which of the below solutions will meet this demand.

- A. Change the Instance type in the current launch configuratio
- B. Change the Desired value of the Autoscaling Group to 4. Ensure the new instances are launched.
- C. Delete the current Launch configuratio
- D. Create a new launch configuration with the new instance type and add it to the Autoscaling Grou
- E. This will then launch the new instances.
- F. Make a copy the Launch configuratio
- G. Change the instance type in the new launch configuratio
- H. Attach that to the Autoscaling Group.Change the maximum and Desired size of the Autoscaling Group to 4. Once the new instances are launched, change the Desired and maximum size back to 2.
- I. Change the desired and maximum size of the Autoscaling Group to 4. Make a copy the Launch configuratio
- J. Change the instance type in the new launch configuratio
- K. Attach that to the Autoscaling Grou
- L. Change the maximum and Desired size of the Autoscaling Group to 2

**Answer: C**

**Explanation:**

You should make a copy of the launch configuration, add the new instance type. The change the Autoscaling Group to include the new instance type. Then change the Desired number of the Autoscaling Group to 4 so that instances of new instance type can be launched. Once launched, change the desired size back to 2, so that Autoscaling will delete the instances with the older configuration. Note that the assumption here is that the current instances are equally distributed across multiple AZ's because Autoscaling will first use the AZRebalance process to terminate instances.

Option A is invalid because you cannot make changes to an existing Launch configuration.

Option B is invalid because if you delete the existing launch configuration, then your application will not be available. You need to ensure a smooth deployment process.

Option D is invalid because you should change the desired size to 4 after attaching the new launch configuration.

For more information on Autoscaling Suspend and Resume, please visit the below URL: <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resume-processes.html>

**NEW QUESTION 189**

You currently have an Autoscalinggroup that has the following settings Min capacity-2

Desired capacity - 2 Maximum capacity - 4

The current number of instances running in the Autoscaling Group is 2. You have been notified that for a duration of an hour, you need to ensure that no new instances are launched by the Autoscaling Group Which of the below 2 actions can be carried out to fulfil this requirement

- A. Change the Maximum capacity to 2
- B. Change the Desired capacity to 4
- C. Suspend the Launch process of the Autoscaling Group
- D. Change the Minimum capacity to 2

**Answer: AC**

**Explanation:**

You can temporarily suspend the creation of new instances by either reducing the Maximum capacity to 2, so that the current instances running which is 2, matches the maximum limit.

Secondly you can suspend the launch process of the Autoscaling Group The AWS Documentation mentions Scaling Processes

Amazon EC2 Auto Scaling supports the following scaling processes:

Launch Adds a new EC2 instance to the group, increasing its capacity. Warning

If you suspend Launch, this disrupts other processes. For example, you can't return an instance in a standby state to service if the Launch process is suspended, because the group can't scale. For more information on Autoscaling Suspend and Resume, please visit the below URL:

? <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resume-processes.html>

**NEW QUESTION 191**

An audit is going to be conducted for your company's AWS account. Which of the following steps will ensure that the auditor has the right access to the logs of your AWS account

- A. Enable S3 and ELB log
- B. Send the logs as a zip file to the IT Auditor.
- C. Ensure CloudTrail is enable
- D. Create a user account for the Auditor and attach the AWSCloudTrailReadOnlyAccess Policy to the user.
- E. Ensure that Cloudtrail is enable
- F. Create a user for the IT Auditor and ensure that full control is given to the userfor Cloudtrail.D- Enable Cloudwatch log
- G. Create a user for the IT Auditor and ensure that full control is given to the userfor the Cloudwatch logs.

**Answer: B**

**Explanation:**

The screenshot shows the 'User details' section for a user named 'testuser'. The details include:

- User name: testuser
- AWS access type: AWS Management Console access - with a password
- Console password type: Autogenerated
- Require password reset: Yes
- Permissions boundary: Permissions boundary is not set

The 'Permissions summary' section shows the following policies attached to the user:

Type	Name
Managed policy	AWSCloudTrailReadOnlyAccess
Managed policy	IAMUserChangePassword

The AWS Documentation clearly mentions the below

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

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

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

**NEW QUESTION 194**

Your company has a set of EC2 resources hosted on AWS. Your new IT procedures state that AWS EC2 Instances must be of a particular Instance type. Which of the following can be used to get the list of EC2 Instances which currently don't match the instance type specified in the new IT procedures

- A. Use AWS Cloudwatch alarms to check which EC2 Instances don't match the intended instance type.
- B. Use AWS Config to create a rule to check the EC2 Instance type
- C. Use Trusted Advisor to check which EC2 Instances don't match the intended instance type.
- D. Use VPC Flow Logs to check which EC2 Instances don't match the intended instance type.

**Answer: B**

**Explanation:**

In AWS Config, you can create a rule which can be used to check if EC2 Instances follow a particular instance type. Below is a snapshot of the output of a rule to check if EC2 instances matches the type of t2micro.

The screenshot shows the configuration for an AWS Config rule. The description is "Checks whether your EC2 instances are of the specified instance types." The trigger type is "Configuration changes". The scope of changes is "Resources". The resource types are "EC2 Instance". The config rule ARN is "arn:aws:config:ap-southeast-1:213171387512:config-rule/config-rule-1exnif". The parameters are "instanceType: t2.micro". The overall rule status shows "Last successful invocation on August 29, 2017 at 12:06:57 PM" and "Last successful evaluation on August 29, 2017 at 12:06:59 PM", both with green checkmarks.

**Resources evaluated**

Click on the icon to view configuration details for the resource when it was last evaluated with this rule.

Resource type	Config timeline	Compliance	Last successful invocation	Last successful evaluation	Manage resource
EC2 Instance	i-0798cb2c515b3aa50	Noncompliant	August 29, 2017 11:25:10 AM	August 29, 2017 11:25:12 AM	
EC2 Instance	i-0f2a1120bec8db3d0	Noncompliant	August 29, 2017 12:06:57 PM	August 29, 2017 12:06:59 PM	

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

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

**NEW QUESTION 198**

Your company is planning to setup a wordpress application. The wordpress application will connect to a MySQL database. Part of the requirement is to ensure that the database environment is fault tolerant and highly available. Which of the following 2 options individually can help fulfil this requirement.

- A. Create a MySQL RDS environment with Multi-AZ feature enabled
- B. Create a MySQL RDS environment and create a Read Replica
- C. Create multiple EC2 instances in the same A
- D. Host MySQL and enable replication via scripts between the instances.
- E. Create multiple EC2 instances in separate AZ'
- F. Host MySQL and enable replication via scripts between the instances.

**Answer: AD**

**Explanation:**

One way to ensure high availability and fault tolerant environments is to ensure Instances are located across multiple availability zones. Hence if you are hosting MySQL yourself, ensure you have instances spread across multiple AZ's

The AWS Documentation mentions the following about the multi-AZ feature

Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. Amazon RDS uses several different technologies to provide failover support. Multi-AZ deployments for Oracle, PostgreSQL, MySQL, and MariaDB DB instances use Amazon's failover technology

For more information on AWS Multi-AZ deployments, please visit the below URL:

- <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html>

**NEW QUESTION 200**

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

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

**Answer:** C

**Explanation:**

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

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

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

**NEW QUESTION 201**

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

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

**Answer:** C

**Explanation:**

Answer - C

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

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

[http://docs.aws.amazon.com/IAM/latest/UserGuide/access\\_policies.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html)

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

The AWS documentation mentions the following

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

For more information on web identity federation please refer to the below link: [http://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles\\_providers\\_oidc.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html)

**NEW QUESTION 205**

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

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

**Answer:** A

**Explanation:**

The AWS Documentation mentions

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

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

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

**NEW QUESTION 208**

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

\* One year free update

You can enjoy free update one year. 24x7 online support.

\* Trusted by Millions

We currently serve more than 30,000,000 customers.

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your DOP-C01 Exam with Our Prep Materials Via below:**

<https://www.certleader.com/DOP-C01-dumps.html>