
Question: 1

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

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

Answer: B

Explanation:

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

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

For more information on On-demand scaling, please refer to the below link:

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

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

Question: 2

You have decided that you need to change the instance type of your production instances which are running as part of an AutoScaling group. The entire architecture is deployed using CloudFormation Template. You currently have 4 instances in Production. You cannot have any interruption in service and need to ensure 2 instances are always running during the update? Which of the options below listed can be used for this?

- A. AutoScalingRollingUpdate
- B. AutoScalingScheduledAction
- C. AutoScalingReplacingUpdate
- D. AutoScalingIntegrationUpdate

Answer: A

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. For more information on Autoscaling updates, please refer to the below link:

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

Question: 3

You currently have the following setup in AWS

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

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

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

Answer: A

Explanation:

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

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

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

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

For more information on Route53 routing policies, please refer to the below link:

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

Question: 4

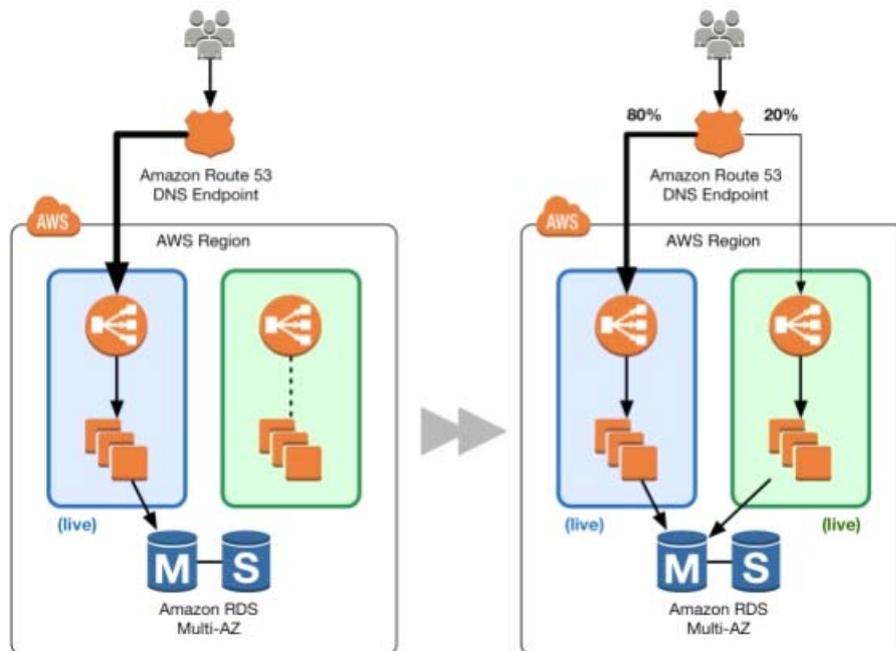
Your application is currently running on Amazon EC2 instances behind a load balancer. Your management has decided to use a Blue/Green deployment strategy. How should you implement this for each deployment?

- 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 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%20Blue%20Green%20Deployments.pdf)

Question: 5

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>

Question: 6

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>

Question: 7

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>

Question: 8

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

Answer: C, D

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 CBS 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>

Question: 9

Which Auto Scaling process would be helpful when testing new instances before sending traffic to them, while still keeping them in your Auto Scaling Group?

- A. Suspend the process AZ Rebalance
- B. Suspend the process Health Check
- C. Suspend the process Replace Unhealthy
- D. Suspend the process AddToLoadBalancer

Answer: D

Explanation:

If you suspend AddTo Load Balancer, Auto Scaling launches the instances but does not add them to the load balancer or target group. If you resume the AddTo Load Balancer process, Auto Scaling resumes adding instances to the load balancer or target group when they are launched. However, Auto Scaling does not add the instances that were launched while this process was suspended. You must register those instances manually.

Option A is invalid because this just balances the number of EC2 instances in the group across the Availability Zones in the region

Option B is invalid because this just checks the health of the instances. Auto Scaling marks an instance as unhealthy if Amazon EC2 or Elastic Load Balancing tells

Auto Scaling that the instance is unhealthy.

Option C is invalid because this process just terminates instances that are marked as unhealthy and later creates new instances to replace them.

For more information on process suspension, please refer to the below document link: from AWS <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resume-processes.html>

Question: 10

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 ELB 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>

Question: 11

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. 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. 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. Workers then use Amazon Simple Email Service to send messages to your on call teams.

B. Create an Amazon SNS topic and configure your on-call team email addresses as subscribers. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to this new topic. Notifications will be sent to on-call users when a CloudWatch alarm is triggered.

C. Create an Amazon SNS topic and configure your on-call team email addresses as subscribers. Create a secondary Amazon SNS topic for alarms and configure your CloudWatch alarms to notify this topic when triggered. Create an HTTP subscriber to this topic that notifies your application via HTTP POST when an alarm is triggered. 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.

D. Create an Amazon SNS topic for each on-call group, and configure each of these with the team member emails as subscribers. Create another Amazon SNS topic and configure your CloudWatch alarms to notify this topic when triggered. Create an HTTP subscriber to this topic that notifies your application via HTTP POST when an alarm is triggered. 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 fulfils all the requirements

- 1) First is to create a SNS topic for each group so that the required members get the email addresses.
- 2) Ensure the application uses the HTTPS endpoint and the SDK to publish messages Option A is invalid because the SQS service is not required.

Option B and C are incorrect. As per the requirement we need to provide notification to only those on-call teams who are working in that particular shift when an alarm is triggered. It need not have to be send to all the on-call teams of the company. With Option B & C, since we are not configuring the SNS topic for each on call team the notifications will be send to all the on-call teams. Hence these 2 options are invalid. For more information on setting up notifications, please refer to the below document link: from AWS

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

Question: 12

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. Configure Elastic Load Balancing health checks to initiate a new AWS CloudFormation stack when health checks return failed.
- B. Re-deploy your infrastructure using an AWS CloudFormation template. Spin up a second AWS CloudFormation stack. Configure Elastic Load Balancing SpillOver functionality to spill over any slow connections to the second AWS CloudFormation stack.
- C. Re-deploy your infrastructure using AWS CloudFormation, Elastic Beanstalk, and Auto Scaling. Set up your Auto Scaling group policies to scale based on the number of requests per second as well as the current customer load time. •>/
- D- Re-deploy your application using an Auto Scaling template. 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>

Question: 13

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

- A. Push custom metrics to CloudWatch to monitor your CPU and network bandwidth from your servers, which will allow your Auto Scaling policy to have better fine-grain insight.
- B. Increase your Auto Scaling group's number of max servers.
- C. Create a script that runs and monitors your servers; when it detects an anomaly in load, it posts to an Amazon SNS topic that triggers Elastic Load Balancing to add more servers to the load balancer.

D. Push custom metrics to CloudWatch for your application that include more detailed information about your web application, such as how many requests it is handling and how many are waiting to be processed.

Answer: B, D

Explanation:

Option B makes sense because maybe the max servers is low hence the application cannot handle the peak load.

Option D helps in ensuring Autoscaling can scale the group on the right metrics.

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

Question: 14

Management has reported an increase in the monthly bill from Amazon Web Services, and they are extremely concerned with this increased cost. Management has asked you to determine the exact cause of this increase. After reviewing the billing report, you notice an increase in the data transfer cost. How can you provide management with a better insight into data transfer use?

A. Update your Amazon CloudWatch metrics to use five-second granularity, which will give better detailed metrics that can be combined with your billing data to pinpoint anomalies.

B. Use Amazon CloudWatch Logs to run a map-reduce on your logs to determine high usage and data transfer.

C. Deliver custom metrics to Amazon CloudWatch per application that breaks down application data transfer into multiple, more specific data points.

D- Using Amazon CloudWatch metrics, pull your Elastic Load Balancing outbound data transfer metrics monthly, and include them with your billing report to show which application is causing higher bandwidth usage.

Answer: C,

Explanation:

You can publish your own metrics to CloudWatch using the AWS CLI or an API. You can view statistical graphs of your published metrics with the AWS Management Console.

CloudWatch stores data about a metric as a series of data points. Each data point has an associated time stamp. You can even publish an aggregated set of data points called a statistic set.

If you have custom metrics specific to your application, you can give a breakdown to the management on the exact issue.

Option A won't be sufficient to provide better insights.

Option B is an overhead when you can make the application publish custom metrics

Option D is invalid because just the ELB metrics will not give the entire picture

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

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

Question: 15

You have an application consisting of a stateless web server tier running on Amazon EC2 instances behind load balancer, and are using Amazon RDS with read replicas. Which of the following methods should you use to implement a self-healing and cost-effective architecture? Choose 2 answers from the options given below

- A. Set up a third-party monitoring solution on a cluster of Amazon EC2 instances in order to emit custom Cloud Watch metrics to trigger the termination of unhealthy Amazon EC2 instances.
- B. Set up scripts on each Amazon EC2 instance to frequently send ICMP pings to the load balancer in order to determine which instance is unhealthy and replace it.
- C. Set up an Auto Scaling group for the web server tier along with an Auto Scaling policy that uses the Amazon RDS DB CPU utilization Cloud Watch metric to scale the instances.
- D. Set up an Auto Scaling group for the web server tier along with an Auto Scaling policy that uses the Amazon EC2 CPU utilization CloudWatch metric to scale the instances.
- E. Use a larger Amazon EC2 instance type for the web server tier and a larger DB instance type for the data storage layer to ensure that they don't become unhealthy.
- F. Set up an Auto Scaling group for the database tier along with an Auto Scaling policy that uses the Amazon RDS read replica lag CloudWatch metric to scale out the Amazon RDS read replicas.
- G. Use an Amazon RDS Multi-AZ deployment.

Answer: D, G

Explanation:

The scaling of EC2 Instances in the Autoscaling group is normally done with the metric of the CPU utilization of the current instances in the Autoscaling group

For more information on scaling in your Autoscaling Group, please refer to the below link:

- <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scaling-simple-step.html>

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. For more information on RDS Multi-AZ please refer to the below link:

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

Option A is invalid because if you already have in-built metrics from Cloudwatch, why would you want to spend more in using a third-party monitoring solution.

Option B is invalid because health checks are already a feature of AWS CLB

Option C is invalid because the database CPU usage should not be used to scale the web tier.

Option E is invalid because increasing the instance size does not always guarantee that the solution will not become unhealthy.

Option F is invalid because increasing Read-Replica's will not suffice for write operations if the primary DB fails.

Question: 16

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

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

Answer: C,

Explanation:

You can use Cloud Watch Logs to monitor applications and systems using log data. For example, CloudWatch Logs can track the number of errors that occur in your application logs and send you a notification whenever the rate of errors exceeds a threshold you specify. CloudWatch Logs uses your log data for monitoring; so, no code changes are required.

Option A and B are invalid because Autoscaling policies are not designed for these purposes.

Option D is invalid because you use Cloudwatch Logs Agent and not the monitoring agent.

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

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

Question: 17

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

- A. Add an Amazon S3 bucket policy with a condition statement to allow access only from Amazon EC2 instances with RFC 1918 IP addresses and enable bucket versioning.

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

Answer: B, D

Explanation:

You can add another layer of protection by enabling MFA Delete on a versioned bucket. Once you do so, you must provide your AWS account's access keys and a valid code from the account's MFA device in order to permanently delete an object version or suspend or reactivate versioning on the bucket.

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

<https://aws.amazon.com/blogs/security/securing-access-to-aws-using-mfa-part-3/>

IAM roles are designed so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that the applications use. Instead of creating and distributing your AWS credentials, you can delegate permission to make API requests using IAM roles. For more information on Roles for EC2 please refer to the below link:

<http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

Option A is invalid because this will not address either the integrity or security concern completely.

Option C is invalid because user credentials should never be used in EC2 instances to access AWS resources.

Option E and F are invalid because AWS Pipeline is an unnecessary overhead when you already have inbuilt controls to manage security for S3.

Question: 18

The project you are working on currently uses a single AWS CloudFormation template to deploy its AWS infrastructure, which supports a multi-tier web application. You have been tasked with organizing the AWS CloudFormation resources so that they can be maintained in the future, and so that different departments such as Networking and Security can review the architecture before it goes to Production. How should you do this in a way that accommodates each department, using their existing workflows?

- A. Organize the AWS CloudFormation template so that related resources are next to each other in the template, such as VPC subnets and routing rules for Networking and security groups and IAM information for Security.

-
- B. Separate the AWS CloudFormation template into a nested structure that has individual templates for the resources that are to be governed by different departments, and use the outputs from the networking and security stacks for the application template that you control. ^/
- C. Organize the AWS CloudFormation template so that related resources are next to each other in the template for each department's use, leverage your existing continuous integration tool to constantly deploy changes from all parties to the Production environment, and then run tests for validation.
- D. Use a custom application and the AWS SDK to replicate the resources defined in the current AWS CloudFormation template, and use the existing code review system to allow other departments to approve changes before altering the application for future deployments.

Answer: B

Explanation:

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

For more information on best practices for CloudFormation please refer to the below link:

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

Question: 19

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: A, C

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:

[Error! Hyperlink reference not valid.](#)

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>

Question: 20

You have the following application to be setup in AWS

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

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

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

Answer: D

Explanation:

IAM roles are designed so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that the applications use. Instead of creating and distributing your AWS credentials

For more information on IAM Roles please refer to the below link:

<http://docs.aws.amazon.com/AWSC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

Question: 21

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: B, C

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/>

Question: 22

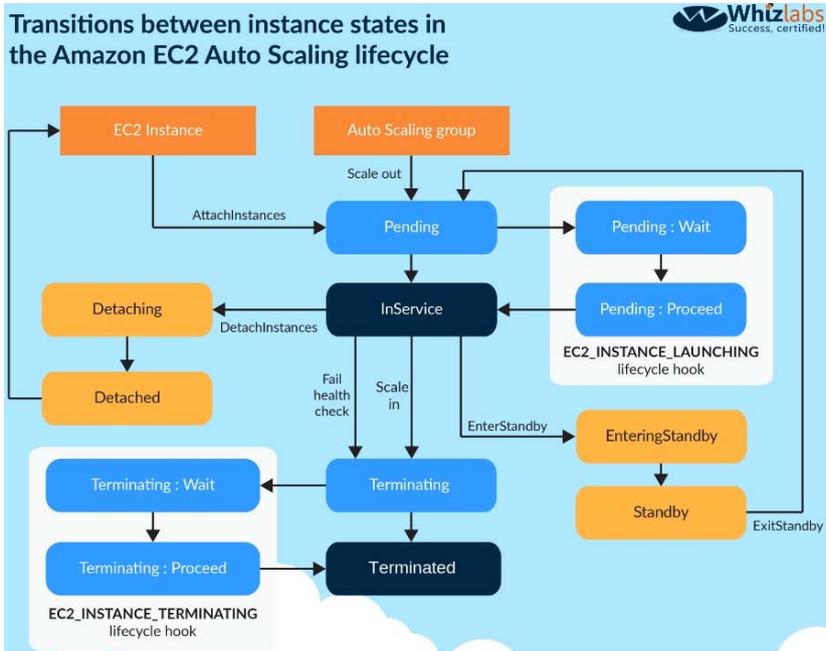
You have an Auto Scaling group of Instances that processes messages from an Amazon Simple Queue Service (SQS) queue. The group scales on the size of the queue. Processing involves calling a third-party web service. The web service is complaining about the number of failed and repeated calls it is receiving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

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

Answer: D

Explanation:

The following diagram shows the lifecycle of the instances in Autoscaling



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

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

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

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>

Question: 23

Your mobile application includes a photo-sharing service that is expecting tens of thousands of users at launch. You will leverage Amazon Simple Storage Service (S3) for storage of the user images, and you must decide how to authenticate and authorize your users for access to these images. You also need to manage the storage of these images. Which two of the following approaches should you use? Choose two answers from the options below

- A. Create an Amazon S3 bucket per user, and use your application to generate the S3 URI for the appropriate content.
- B. Use AWS Identity and Access Management (IAM) user accounts as your application-level user database, and offload the burden of authentication from your application code.
- C. Authenticate your users at the application level, and use AWS Security Token Service (STS) to grant token-based authorization to S3 objects.

- D. Authenticate your users at the application level, and send an SMS token message to the user. Create an Amazon S3 bucket with the same name as the SMS message token, and move the user's objects to that bucket.
- E. Use a key-based naming scheme comprised from the user IDs for all user objects in a single Amazon S3 bucket.

Answer: C, E

Explanation:

The AWS Security Token Service (STS) is a web service that enables you to request temporary, limited-privilege credentials for AWS Identity and Access

Management (IAM) users or for users that you authenticate (federated users).

The token can then be used to grant access to the objects in S3.

You can then provides access to the objects based on the key values generated via the user id.

Option A is possible but then becomes a maintenance overhead because of the number of buckets.

Option B is invalid because IAM users is not a good security practice.

Option D is invalid because SMS tokens are not efficient for this requirement.

For more information on the Security Token Service please refer to the below link:

<http://docs.aws.amazon.com/STS/latest/APIReference/Welcome.html>

Question: 24

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

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

Answer: A

Explanation:

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

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

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

Question: 25

You have an Auto Scaling group with 2 AZs. One AZ has 4 EC2 instances and the other has 3 EC2 instances. None of the instances are protected from scale in. Based on the default Auto Scaling termination policy what will happen?

- A. Auto Scaling selects an instance to terminate randomly
- B. Auto Scaling will terminate unprotected instances in the Availability Zone with the oldest launch configuration.
- C. Auto Scaling terminates which unprotected instances are closest to the next billing hour.
- D. Auto Scaling will select the AZ with 4 EC2 instances and terminate an instance.

Answer: D

Explanation:

The default termination policy is designed to help ensure that your network architecture spans Availability Zones evenly. When using the default termination policy.

Auto Scaling selects an instance to terminate as follows:

Auto Scaling determines whether there are instances in multiple Availability Zones. If so, it selects the Availability Zone with the most instances and at least one instance that is not protected from scale in. If there is more than one Availability Zone with this number of instances. Auto Scaling selects the Availability Zone with the instances that use the oldest launch configuration. For more information on Autoscaling instance termination please refer to the below link:

<http://docs.aws.amazon.com/autoscaling/latest/userguide/as-instance-termination.html>

Question: 26

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 Scalinggroup 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: A, D, E

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 OastiCache 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 OastiCache is an in-memory key/value store that sits between ycbetappiicipiJGra arcdalie data store (database) that it accesses. Whenever your application requests data, it first makes the request to the DastiCache cache. If the data exists in the cache and is current, OastiCache 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 data store which returns the data to your application. Your application then writes the data received from the store to the cache so it can be more quickly retrieved next time it is requested. For more information on Elastic Cache 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

Question: 27

You are administering a continuous integration application that polls version control for changes and then launches new Amazon EC2 instances for a full suite of build tests. What should you do to ensure the lowest overall cost while being able to run as many tests in parallel as possible?

A. Perform syntax checking on the continuous integration system before launching a new Amazon EC2 instance for build test, unit and integration tests.

B. Perform syntax and build tests on the continuous integration system before launching the new Amazon EC2 instance unit and integration tests. C. Perform all tests on the continuous integration system, using AWS OpsWorks for unit, integration, and build tests.

D. Perform syntax checking on the continuous integration system before launching a new AWS Data Pipeline for coordinating the output of unit, integration, and build tests.

Answer: B

Explanation:

Continuous Integration (CI) is a development practice that requires developers to integrate code into a shared repository several times a day. Each check-in is then verified by an automated build, allowing teams to detect problems early.

Option A and D are invalid because you can do build tests on a CI system and not only Syntax tests. And Syntax tests are normally done during coding time and not during the build time.

Option C is invalid because Opswork is ideally not used for build and integration tests.

For an example of a Continuous integration system, please refer to the Jenkins system via the url below

- <https://jenkins.io/>

Question: 28

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

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

Answer: A

Explanation:

Elastic Beanstalk simplifies this process by managing the Amazon SQS queue and running a daemon process on each instance that reads from the queue for you.

When the daemon pulls an item from the queue, it sends an HTTP POST request locally to `http://localhost/` with the contents of the queue message in the body. All that your application needs to do is perform the long-running task in response to the POST.

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

- ◆ <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.html>

Question: 29

You have an Auto Scaling group with an Elastic Load Balancer. You decide to suspend the Auto Scaling AddToLoadBalancer for a short period of time. What will happen to the instances launched during the suspension period?

- A. The instances will be registered with ELB once the process has resumed
- B. Auto Scaling will not launch the instances during this period because of the suspension
- C. The instances will not be registered with ELB. You must manually register when the process is resumed */
- D. It is not possible to suspend the AddToLoadBalancer process

Answer: C

Explanation:

If you suspend AddTo Load Balancer, Auto Scaling launches the instances but does not add them to the load balancer or target group. If you resume the AddTo Load Balancer process, Auto Scaling resumes adding instances to the load balancer or target group when they are launched. However, Auto Scaling does not add the instances that were launched while this process was suspended. You must register those instances manually.

For more information on the Suspension and Resumption process, please visit the below U RL:

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

Question: 30

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. Only execute the script if "true" is returned.
- C. Use a "Solo Command" within an Elastic Beanstalk configuration file to execute the script. The Elastic Beanstalk service will ensure that the command is only executed once.
- D. 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>