

Exam Questions DOP-C01

AWS Certified DevOps Engineer- Professional

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



NEW QUESTION 1

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/>

NEW QUESTION 2

You currently have the following setup in AWS

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

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

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

Answer: A

Explanation:

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

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

Option C is wrong because 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>

NEW QUESTION 3

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 AddToLoadBalancer, Auto Scaling launches the instances but does not add them to the load balancer or target group. If you resume the AddToLoadBalancer 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>

NEW QUESTION 4

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

- A. Create an Amazon SNS topic and an Amazon SQS queue

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

Answer: D

Explanation:

Option D fulfills all the requirements

1) First is to create a SNS topic for each group so that the required members get the email addresses.

2) Ensure the application uses the HTTPS endpoint and the SDK to publish messages Option A is invalid because the SQS service is not required.

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

NEW QUESTION 5

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 6

You have an application consisting of a stateless web server tier running on Amazon EC2 instances behind a 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: DG

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 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 C 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.

NEW QUESTION 7

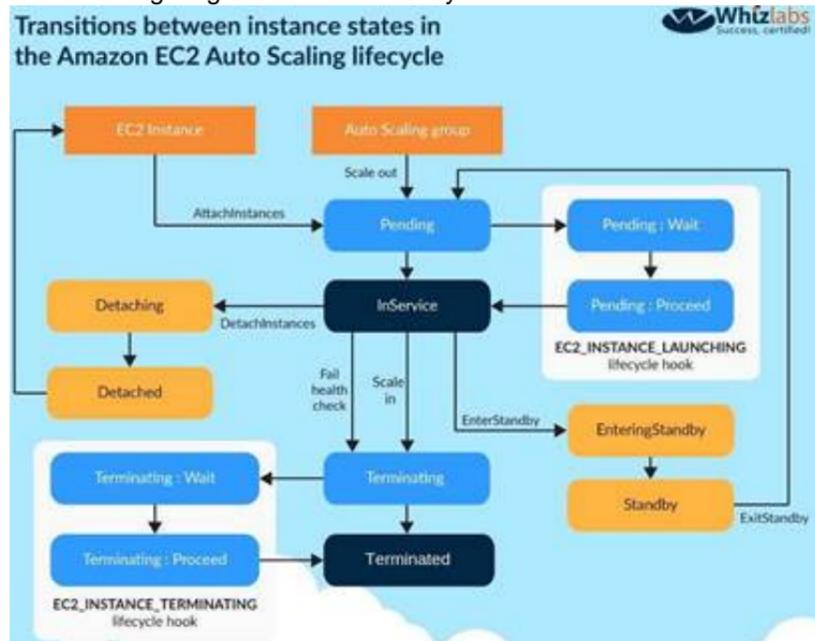
You have an Auto Scaling group of Instances that processes messages from an Amazon Simple Queue Service (SQS) queue. The group scales on the size of the queue. Processing Involves calling a third- party web service. The web service is complaining about the number of failed and repeated calls it is receiving 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 softwar
- 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 Scalinggroup, and change the scaling policies so they scale less dynamically.
- E. Modify the application running on the instances to put itself into an Auto Scaling Standby state while it processes a task and return itself to InService when the processing is complete.

Answer: D

Explanation:

The following diagram shows the lifecycle of the instances in Autoscaling



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

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

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

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

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

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

NEW QUESTION 8

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 metricbased 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 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

NEW QUESTION 9

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 10

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 EL
- D. You must manually register when the process is resumed */
- E. 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>

NEW QUESTION 10

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 Scalinggroup to scale your post-processing application which is configured to pull down your log files stored an Amazon S3.
- D. Post your log data to an Amazon Kinesis data stream, and subscribe your log-processing application so that is configured to process your logging data.
- E. Create a multi-AZ Amazon RDS MySQL cluster, post the logging data to MySQL, and run a map reduce job to retrieve the required information on user counts.

Answer: C

Explanation:

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

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

Amazon

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

• <https://aws.amazon.com/kinesis>

NEW QUESTION 12

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

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

Answer: C

Explanation:

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

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

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

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

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

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

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

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

NEW QUESTION 14

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

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

Answer: C

Explanation:

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

Option B is partially correct where you store the data in DynamoDB, but then the number of GET requests would still be high if the entire DynamoDB table had to be traversed and each object compared and updated in S3.

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

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

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

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

NEW QUESTION 17

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

- A. Manually create one template to encompass all the resources that you need for the system, so you only have a single template to version-control.
- B. Create multiple separate templates for each logical part of the system, create nested stacks in AWS CloudFormation, and maintain several templates to version-contro
- C. >/
- D. Create multiple separate templates for each logical part of the system, and provide the outputs from one to the next using an Amazon Elastic Compute Cloud (EC2) instance running the SDK forfinergranularity of control.
- E. Manually construct the networking layer using Amazon Virtual Private Cloud (VPC) because this does not change often, and then use AWS CloudFormation to define all other ephemeral resources.

Answer: B

Explanation:

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

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

NEW QUESTION 18

You use Amazon Cloud Watch as your primary monitoring system for your web application. After a

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

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

Answer: BDE

Explanation:

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

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

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

Amazon CloudWatch uses Amazon SNS to send email. First, create and subscribe to an SNS topic.

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

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

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

NEW QUESTION 21

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. Its 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 Autoscaling health checks, please refer to the below document link: from AWS

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

NEW QUESTION 22

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

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

Answer: C

Explanation:

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

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

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

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

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

NEW QUESTION 26

You have an application running on Amazon EC2 in an Auto Scaling group. Instances are being bootstrapped dynamically, and the bootstrapping takes over 15 minutes to complete. You find that instances are reported by Auto Scaling as being In Service before bootstrapping has completed. You are receiving application alarms related to new instances before they have completed bootstrapping, which is causing confusion. You find the cause: your application monitoring tool is polling the Auto Scaling Service API for instances that are In Service, and creating alarms for new previously unknown instances. Which of the following will ensure that new instances are not added to your application monitoring tool before bootstrapping is completed?

- A. Create an Auto Scaling group lifecycle hook to hold the instance in a pending: wait state until your bootstrapping is complete
- B. Once bootstrapping is complete, notify Auto Scaling to complete the lifecycle hook and move the instance into a pending: proceed state.
- C. Use the default Amazon Cloud Watch application metrics to monitor your application's health

- D. Configure an Amazon SNS topic to send these Cloud Watch alarms to the correct recipients.
- E. Tag all instances on launch to identify that they are in a pending state
- F. Change your application monitoring tool to look for this tag before adding new instances, and then use the Amazon API to set the instance state to 'pending' until bootstrapping is complete.
- G. Increase the desired number of instances in your Auto Scaling group configuration to reduce the time it takes to bootstrap future instances.

Answer: A

Explanation:

Auto Scaling lifecycle hooks enable you to perform custom actions as Auto Scaling launches or terminates instances. For example, you could install or configure software on newly launched instances, or download log files from an instance before it terminates. After you add lifecycle hooks to your Auto Scaling group, they work as follows:

1. Auto Scaling responds to scale out events by launching instances and scale in events by terminating instances.
2. Auto Scaling puts the instance into a wait state (Pending:Wait or Terminating:Wait). The instance remains in this state until either you tell Auto Scaling to continue or the timeout period ends.

For more information on rolling updates, please visit the below link:

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

NEW QUESTION 28

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

Answer: B

Explanation:

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

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

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

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 32

When thinking of AWS Elastic Beanstalk's model, which is true?

- A. Applications have many deployments, deployments have many environments.
- B. Environments have many applications, applications have many deployments.
- C. Applications have many environments, environments have many deployments.
- D. Deployments have many environments, environments have many applications.

Answer: C

Explanation:

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

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

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

NEW QUESTION 36

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 instance
- B. Use an S3 Notification Configuration on the PutBucket event and publish events to AWS Lambda
- C. Use the Lambda to analyze logs as soon as they come in and flag issues.
- D. Begin using CloudWatch Logs on every service
- 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 39

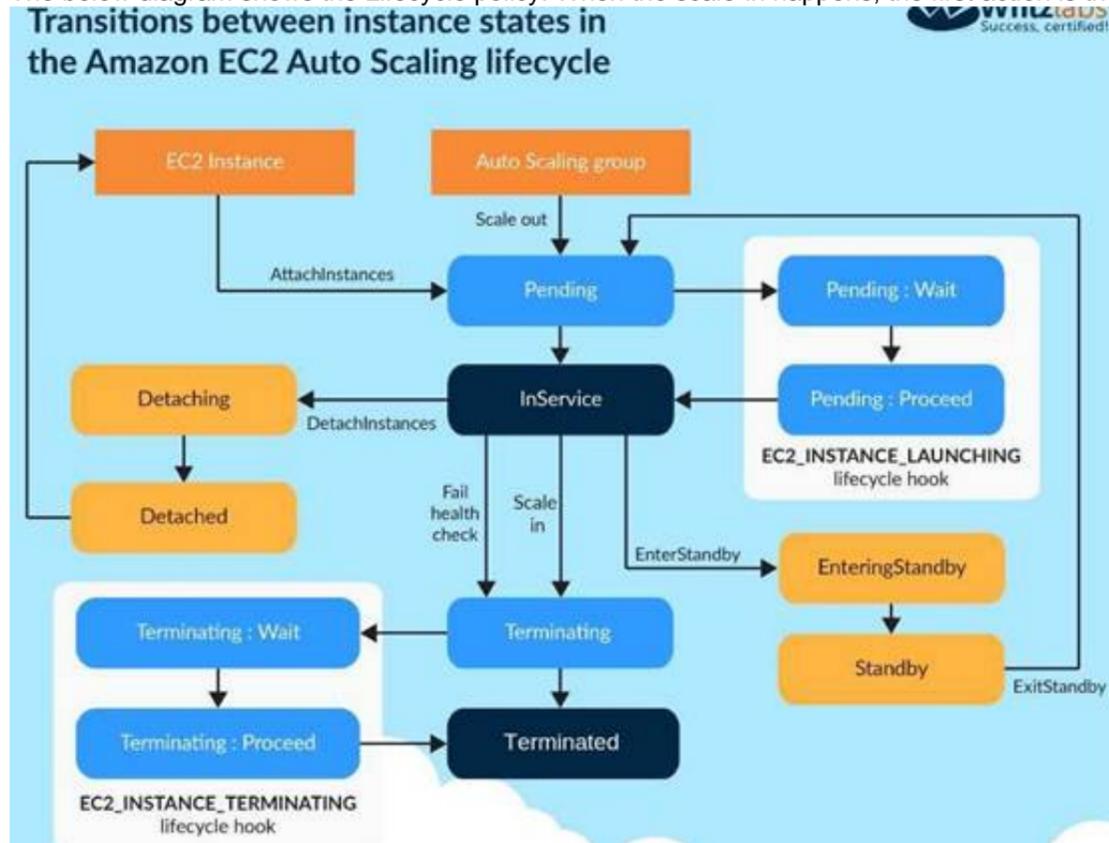
For AWS Auto Scaling, what is the first transition state an instance enters after leaving steady state when scaling in due to health check failure or decreased load?

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

Answer: A

Explanation:

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



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

NEW QUESTION 40

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

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 44

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

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

Answer: C

Explanation:

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

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

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

Option D is invalid because there is a limit of 5 VPCs per region. For more information on Elastic IP's, please visit the below URL:

- <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html>

NEW QUESTION 49

You need to create a simple, holistic check for your system's general availability and uptime. Your system presents itself as an HTTP-speaking API. What is the most simple tool on AWS to achieve this with?

- A. Route53 Health Checks
- B. CloudWatch Health Checks
- C. AWS ELB Health Checks
- D. EC2 Health Checks

Answer: A

Explanation:

Amazon Route 53 health checks monitor the health and performance of your web applications, web servers, and other resources. Each health check that you create

can monitor one of the following:

- The health of a specified resource, such as a web server
- The status of an Amazon Cloud Watch alarm
- The status of other health checks

For more information on Route53 Health checks, please refer to the below link:

- <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html>

NEW QUESTION 54

You need to perform ad-hoc business analytics queries on well-structured data. Data comes in constantly at a high velocity. Your business intelligence team can understand SQL.

What AWS service(s) should you look to first?

- A. Kinesis Firehose + RDS
- B. Kinesis Firehose+RedShift
- C. EMR using Hive
- D. EMR running Apache Spark

Answer: B

Explanation:

Amazon Kinesis Firehose is the easiest way to load streaming data into AWS. It can capture, transform, and load streaming data into Amazon Kinesis Analytics, Amazon S3, Amazon Redshift, and Amazon Elasticsearch Service, enabling near real-time analytics with existing business intelligence tools and dashboards you're already using today. It is a fully managed service that automatically scales to match the throughput of your data and requires no ongoing administration. It can also batch, compress, and encrypt the data before loading it, minimizing the amount of storage used at the destination and increasing security.

For more information on Kinesis firehose, please visit the below URL:

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

Amazon Redshift is a fully managed, petabyte-scale data warehouse service in the cloud. You can start with just a few hundred gigabytes of data and scale to a petabyte or more. This enables you to use your data to acquire new insights for your business and customers. For more information on Redshift, please visit the below URL:

<http://docs.aws.amazon.com/redshift/latest/mgmt/welcome.html>

NEW QUESTION 56

Your company needs to automate 3 layers of a large cloud deployment. You want to be able to track this deployment's evolution as it changes over time, and carefully control any alterations. What is a good way to automate a stack to meet these requirements?

- A. Use OpsWorks Stacks with three layers to model the layering in your stack.
- B. Use CloudFormation Nested Stack Templates, with three child stacks to represent the three logical layers of your cloud.
- C. Use AWS Config to declare a configuration set that AWS should roll out to your cloud.
- D. Use Elastic Beanstalk Linked Applications, passing the important DNS entries between layers using the metadata interface.

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 nested stacks, please visit the below URL:

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

The query is, how you can automate a stack over the period of time, when changes are required, with out recreating the stack.

The function of Nested Stacks are to reuse Common Template Patterns.

For example, assume that you have a load balancer configuration that you use for most of your stacks. Instead of copying and pasting the same configurations into your templates, you can create a dedicated template for the load balancer. Then, you just use the resource to reference that template from within other templates. Yet another example is if you have a launch configuration with certain specific configuration and you need to change the instance size only in the production environment and to leave it as it is in the development environment.

AWS also recommends that updates to nested stacks are run from the parent stack.

When you apply template changes to update a top-level stack, AWS CloudFormation updates the top-level stack and initiates an update to its nested stacks. AWS Cloud Formation updates the resources of modified nested stacks, but does not update the resources of unmodified nested stacks.

NEW QUESTION 57

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 regio
- B. Create an Auto Scaling Group behind an ELB in each of the two regions for your application layer in which DynamoDB is running i
- 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 tabl
- 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 i
- 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 tabl
- 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 regio
- J. Create a crossregion 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 sent 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 wort: only if you use a failover routing policy.

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

Amazon DynamoDBglobal 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 GlobalTables, please visit the below URL:

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

NEW QUESTION 58

You are building a Ruby on Rails application for internal, non-production use which uses MySQL as a database. You want developers without very much AWS experience to be able to deploy new code with a single command line push. You also want to set this up as simply as possible. Which tool is ideal for this setup?

- A. AWSCloudFormation
- B. AWSOpsWorks
- C. AWS ELB+ EC2 with CLI Push
- D. AWS Elastic Beanstalk

Answer: D

Explanation:

With Elastic Beanstalk, you can quickly deploy and manage applications in the AWS Cloud without worrying about the infrastructure that runs those applications.

AWS Elastic Beanstalk reduces management complexity without restricting choice or control. You simply upload your application, and Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring

Elastic Beanstalk supports applications developed in Java, PHP, .NET, Node.js, Python, and Ruby, as well as different container types for each language.

For more information on Elastic beanstalk, please visit the below URL:

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

NEW QUESTION 61

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

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

Answer: B

Explanation:

The AWS::AutoScaling::AutoScalingGroup resource supports an UpdatePolicy attribute. This is used to define how an Auto Scalinggroup 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 Scalinggroup and replaces old instances with new ones, according to the parameters specified.

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

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

NEW QUESTION 64

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

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

Answer: C

Explanation:

The AWS documentation mentions

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

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

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

NEW QUESTION 67

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

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

Answer: B

Explanation:

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

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

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

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

NEW QUESTION 71

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

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

Answer: B

Explanation:

The AWS Documentation mentions

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

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

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

NEW QUESTION 73

Which of the following environment types are available in the Elastic Beanstalk environment. Choose 2 answers from the options given below.

- A. Single Instance
- B. Multi-Instance
- C. Load Balancing Autoscaling
- D. SQS, Autoscaling

Answer: AC

Explanation:

The AWS Documentation mentions

In Elastic Beanstalk, you can create a load-balancing, autoscaling environment or a single-instance environment. The type of environment that you require depends

on the application that you deploy.

When you go onto the Configuration for your environment, you will be able to see the Environment type from there.

NEW QUESTION 77

You have a video processing application hosted in AWS. The videos 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

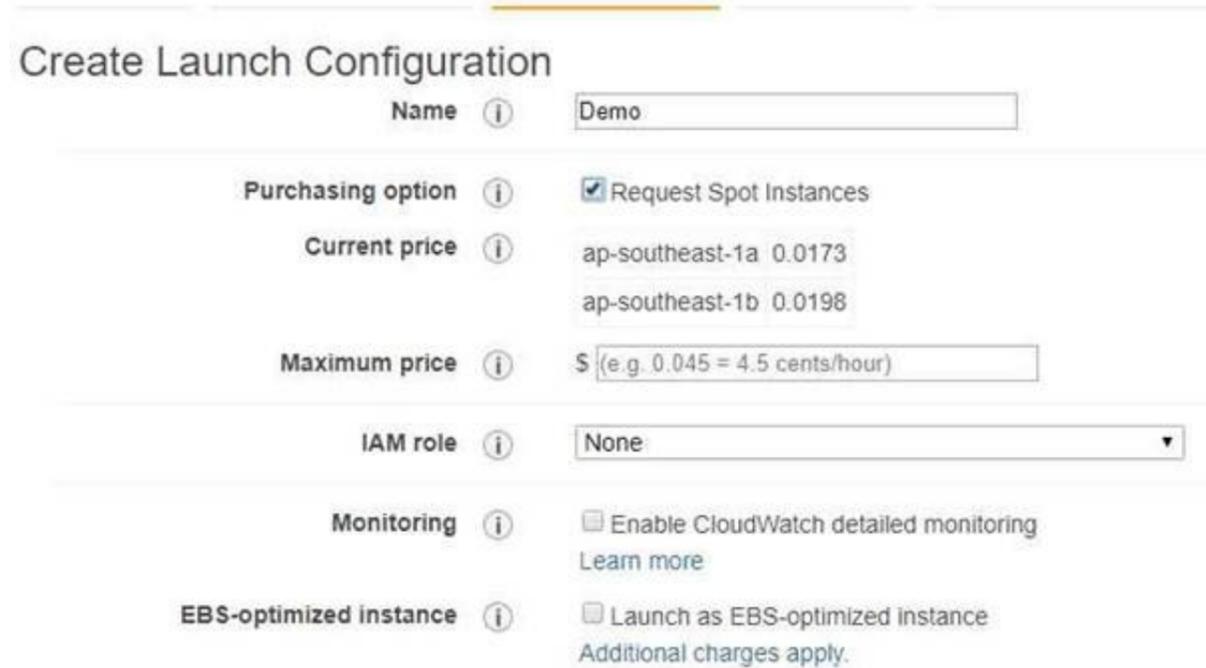
- C. Create an Autoscalinggroup with the launch configuration.
- D. Create a launch configuration with Spot Instance
- E. Ensure the User Data section details the installation of the custom softwar
- F. Create an Autoscalinggroupwith the launch configuration.
- G. Create a launch configuration with Dedicated Instance
- H. Ensure the User Data section details the installation of the custom softwar
- 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 softwar
- 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.



The screenshot shows the 'Create Launch Configuration' page in the AWS console. The 'Name' field is set to 'Demo'. Under 'Purchasing option', the checkbox for 'Request Spot Instances' is checked. The 'Current price' section shows two options: 'ap-southeast-1a 0.0173' and 'ap-southeast-1b 0.0198'. The 'Maximum price' field is set to '\$ (e.g. 0.045 = 4.5 cents/hour)'. The 'IAM role' dropdown is set to 'None'. Under 'Monitoring', the checkbox for 'Enable CloudWatch detailed monitoring' is unchecked. Under 'EBS-optimized instance', the checkbox for 'Launch as EBS-optimized instance' is unchecked, with a note that 'Additional charges apply'.

▶ **Advanced Details**

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

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

NEW QUESTION 81

Your company has the requirement to set up instances running as part of an Autoscaling Group. Part of the requirement is to use Lifecycle hooks to setup custom based software's and do the necessary configuration on the instances. The time required for this setup might take an hour, or might finish before the hour is up. How should you setup lifecycle hooks for the Autoscaling Group. Choose 2 ideal actions you would include as part of the lifecycle hook.

- A. Configure the lifecycle hook to record heartbeat
- B. If the hour is up, restart the timeout period.
- C. Configure the lifecycle hook to record heartbeat
- D. If the hour is up, choose to terminate the current instance and start a new one
- E. If the software installation and configuration is complete, then restart the time period.
- F. If the software installation and configuration is complete, then send a signal to complete the launch of the instance.

Answer: AD

Explanation:

The AWS Documentation provides the following information on lifecycle hooks

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

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

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

NEW QUESTION 83

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 IAM credential and assign the IAM user to a group with proper permissions to communicate with DynamoDB.
- B. Create an IAM 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 IAM role with the proper permission policy to communicate with the DynamoDB table
- E. Use web identity federation, which assumes the IAM 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 IAM 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 IAM policies please refer to the below link:

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

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

The AWS documentation mentions the following

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

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

NEW QUESTION 84

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 88

Your company is planning to develop an application in which the front end is in .Net and the backend is in DynamoDB. There is an expectation of a high load on the application. How could you ensure the scalability of the application to reduce the load on the DynamoDB database? Choose an answer from the options below.

- A. Add more DynamoDB databases to handle the load.
- B. Increase write capacity of Dynamo DB to meet the peak loads
- C. Use SQS to assist and let the application pull messages and then perform the relevant operation in DynamoDB.
- D. Launch DynamoDB in Multi-AZ configuration with a global index to balance writes

Answer: C

Explanation:

When the idea comes for scalability then SQS is the best option. Normally DynamoDB is scalable, but since one is looking for a cost effective solution, the messaging in SQS can assist in managing the situation mentioned in the question.

Amazon Simple Queue Service (SQS) is a fully-managed message queuing service for reliably communicating among distributed software components and microservices - at any scale. Building applications from individual components that each perform a discrete function improves scalability and reliability, and is best practice design for modern applications. SQS makes it simple and cost- effective to decouple and coordinate the components of a cloud application. Using SQS, you can send, store, and receive messages between software components at any volume, without losing messages or requiring other services to be always available

For more information on SQS, please refer to the below URL:

• <https://aws.amazon.com/sqs/>

NEW QUESTION 89

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

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

Answer: B

Explanation:

The AWS documentation mentions

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

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

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

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

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

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

NEW QUESTION 93

Your finance supervisor has set a budget of 2000 USD for the resources in AWS. Which of the following is the simplest way to ensure that you know when this threshold is being reached.

- A. Use Cloudwatch events to notify you when you reach the threshold value
- B. Use the Cloudwatch billing alarm to to notify you when you reach the threshold value
- C. Use Cloudwatch logs to notify you when you reach the threshold value
- D. Use SQS queues to notify you when you reach the threshold value

Answer: B

Explanation:

The AWS documentation mentions

You can monitor your AWS costs by using Cloud Watch. With Cloud Watch, you can create billing alerts that notify you when your usage of your services exceeds thresholds that you define. You specify these threshold amounts when you create the billing alerts.

When your usage exceeds these amounts, AWS sends you an

email notification. You can also sign up to receive notifications when AWS prices change. For more information on billing alarms, please refer to the below URL:

- <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/monitor-charges.html>

NEW QUESTION 94

You are using Autoscaling for managing the instances in your AWS environment. You need to deploy a new version of your application. You'd prefer to use all new instances if possible, but you cannot have any downtime. You also don't want to swap any environment urls. Which of the following deployment methods would you implement

- A. Using "All at once" deployment method.
- B. Using "Blue Green" deployment method.
- C. Using "RollingUpdates" deployment method.
- D. Using "Blue Green" with "All at once" deployment method.

Answer: C

Explanation:

In Rolling deployment, you can mention a new set of servers which can replace the existing set of servers. This replacement will happen in a phased out manner. Since there is a requirement to not swap URL's, you must not use Blue Green deployments.

For more information on the differences between Rolling Updates and Blue Green deployments, please refer to the below URL:

- <https://cloudnative.io/docs/blue-green-deployment/>

NEW QUESTION 97

A user is accessing RDS from an application. The user has enabled the Multi AZ feature with the MS SQL RDS DB. During a planned outage how will AWS ensure that a switch from DB to a standby replica will not affect access to the application?

- A. RDS will have an internal IP which will redirect all requests to the new DB
- B. RDS uses DNS to switch over to stand by replica for seamless transition
- C. The switch over changes Hardware so RDS does not need to worry about access
- D. RDS will have both the DBs running independently and the user has to manually switch over

Answer: B

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 (for example, instance hardware failure, storage failure, or network disruption), Amazon RDS performs an automatic failover to the standby, so that you can resume database operations as soon as the failover is complete.

And as per the AWS documentation, the cname is changed to the standby DB when the primary one fails.

Q: What happens during Multi-AZ failover and how long does it take?

"Failover is automatically handled by Amazon RDS so that you can resume database operations as quickly as possible without administrative intervention. When failing over, Amazon RDS simply flips the canonical name record (CNAMC) for your DB instance to point at the standby, which is in turn promoted to become the new primary. We encourage you to follow best practices and implement database connection retry at the application layer".

<https://aws.amazon.com/rds/faqs/>

Based on this, RDS Multi-AZ will use DNS to create the CNAM C and hence B is the right option. For more information on RDS Multi-AZ please visit the link:

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

NEW QUESTION 101

A user is using Cloudformation to launch an EC2 instance and then configure an application after the instance is launched. The user wants the stack creation of ELB and AutoScaling to wait until the EC2 instance is launched and configured properly. How can the user configure this?

- A. It is not possible that the stack creation will wait until one service is created and launched B. The user can use the HoldCondition resource to wait for the creation of the other dependent resources
- B. The user can use the DependentCondition resource to hold the creation of the other dependent resources
- C. The user can use the WaitCondition resource to hold the creation of the other dependent resources

Answer: D

Explanation:

You can use a wait condition for situations like the following:

To coordinate stack resource creation with configuration actions that are external to the stack creation

To track the status of a configuration process

For more information on Cloudformation Wait condition please visit the link

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

NEW QUESTION 106

You work at a company that makes use of AWS resources. One of the key security policies is to ensure that all data is encrypted both at rest and in transit. Which of the following is not a right implementation which aligns to this policy?

- A. Using S3 Server Side Encryption (SSE) to store the information
- B. Enable SSL termination on the ELB C Enabling Proxy Protocol D- Enabling sticky sessions on your load balancer

Answer: B

Explanation:

Please note the keyword "NOT" in the question.

Option A is incorrect. Enabling S3 SSE encryption helps the encryption of data at rest in S3. So Option A is invalid.

Option B is correct. If you disable SSL termination on the ELB the traffic will be encrypted all the way to the backend. SSL termination allows encrypted traffic between the client

and the ELB but cause traffic to be unencrypted between the ELB and the backend (presumably EC2 or ECS/Task, etc.)

If SSL is not terminated on the ELB you must use Layer A to have traffic encrypted all the way.

Sticky sessions are not supported with Layer A (TCP endpoint). Thus option D "Enabling sticky sessions on your load balancer" can't be used and is the right answer

For more information on sticky sessions, please visit the below URL <https://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-sticky-sessions.html>

Requirements

- An HTTP/HTTPS load balancer.

- At least one healthy instance in each Availability Zone.

- At least one healthy instance in each Availability Zone.

If you don't want the load balancer to handle the SSL termination (known as SSL offloading), you can use TCP for both the front-end and back-end connections, and deploy certificates on the registered instances handling requests.

For more information on elb-listener-config, please visit the below

- <https://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-listener-config.html> If the front-end connection uses TCP or SSL, then your back-end connections can use either TCP or SSL. Note: You can use an HTTPS listener and still use SSL on the backend but the ELB must terminate, decrypt and re-encrypt. This is slower and less secure than using the same encryption all the way to the backend. It also breaks the question requirement of having all data encrypted in transit since it forces the ELB to decrypt. Proxy protocol is used to provide a secure transport connection hence Option C is also incorrect. For more information on SSL Listeners for your load balancer, please visit the below URL

<http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-https-load-balancers.html>

<https://aws.amazon.com/blogs/aws/elastic-load-balancer-support-for-ssl-termination/>

NEW QUESTION 109

Your company is using an Autoscaling Group to scale out and scale in instances. There is an expectation of a peak in traffic every Monday at 8am. The traffic is then expected to come down before the weekend on Friday 5pm. How should you configure Autoscaling in this?

- A. Create dynamic scaling policies to scale up on Monday and scale down on Friday
- B. Create a scheduled policy to scale up on Friday and scale down on Monday
- C. Create a scheduled policy to scale up on Monday and scale down on Friday
- D. Manually add instances to the Autoscaling Group on Monday and remove them on Friday

Answer: C

Explanation:

The AWS Documentation mentions the following for Scheduled scaling

Scaling based on a schedule allows you to scale your application in response to predictable load changes. For example, every week the traffic to your web application starts to increase on Wednesday, remains high on Thursday, and starts to decrease on Friday. You can plan your scaling activities based on the predictable traffic patterns of your web application.

For more information on scheduled scaling for Autoscaling, please visit the below URL

- http://docs.aws.amazon.com/autoscaling/latest/userguide/schedule_time.html

NEW QUESTION 110

You are working for a company that has an on-premise infrastructure. There is now a decision to move to AWS. The plan is to move the development environment first. There are a lot of custom-based applications that need to be deployed for the development community. Which of the following can help to implement the application for the development team?

Choose 2 answers from the options below.

- A. Create Docker containers for the custom application components.
- B. Use OpsWorks to deploy the Docker containers.
- C. Use Elastic Beanstalk to deploy the Docker containers.
- D. Use CloudFormation to deploy the Docker containers.

Answer: AC

Explanation:

The AWS documentation states the following for Docker containers on Elastic Beanstalk

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.

For more information on Docker containers and Elastic Beanstalk, please visit the below URL

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

NEW QUESTION 112

Your firm has uploaded a large amount of aerial image data to S3. In the past, in your on-premise environment, you used a dedicated group of servers to process this data and used RabbitMQ - An open source messaging system to get job information to the servers. Once processed the data would go to tape and be

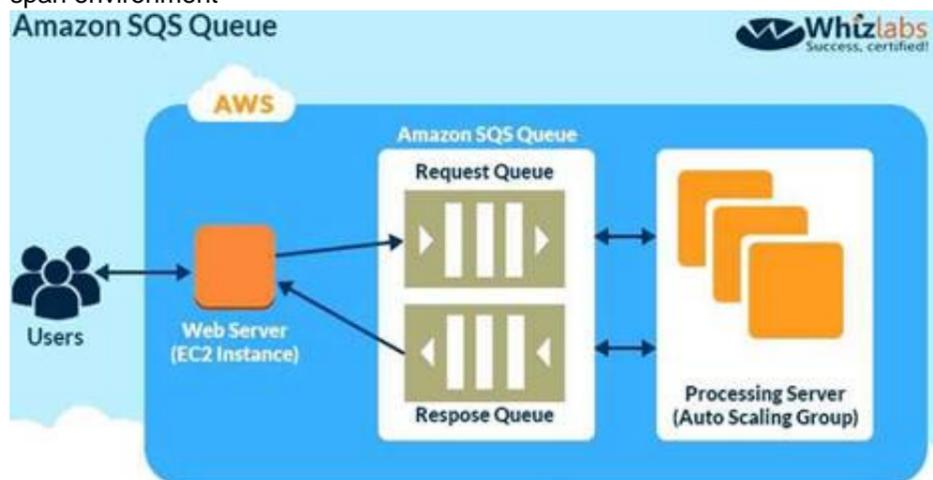
shipped offsite. Your manager told you to stay with the current design, and leverage AWS archival storage and messaging services to minimize cost. Which is correct?

- A. UseSQS for passing job message
- B. Use Cloud Watch alarms to terminate EC2 workerinstances when they become idl
- C. Once data is processed, change the storageclass of the S3 objects to Reduced Redundancy Storage.
- D. SetupAuto-Scaled workers triggered by queue depth that use spot instances to processmessages in SQ
- E. Once data is processed, change the storage class of the S3objects to Glacier
- F. Changethe storage class of the S3 objects to Reduced Redundancy Storag
- G. SetupAuto-Scaled workers triggered by queue depth that use spot instances to processmessages in SQ
- H. Once data is processed, change the storage class of the S3objects to Glacier.
- I. Use SNS to passjob messages use Cloud Watch alarms to terminate spot worker instanceswhen they become idl
- J. Once data is processed, change the storage class of theS3 object to Glacier.

Answer: B

Explanation:

The best option for reduces costs is Glacier, since anyway in the on-premise location everything was stored on tape. Hence option A is out. Next SQS should be used, since RabbitMG was used internally. Hence option D is out. The first step is to leave the objects in S3 and not tamper with that. Hence option B is more suited. The following diagram shows how SQS is used in a worker span environment



For more information on SQS queues, please visit the below URL
<http://docs.ws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-how-it-works.html>

NEW QUESTION 114

You are currently planning on using Autoscaling to launch instances which have an application installed. Which of the following methods will help ensure the instances are up and running in the shortest span of time to take in traffic from the users?

- A. Loginto each instance and install the software.
- B. UseUserData to launch scripts to install the software.
- C. UseDocker containers to launch the software.
- D. UseAMI's which already have the software installed.

Answer: D

Explanation:

The AM I will be the fastest because it will already have the software installed. You can customize the instance that you launch from a public AMI and then save that configuration as a custom AMI for your own use. Instances that you launch from your AMI use all the custom izations that you've made. For more information on AMI'S please refer to the below link <http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/AMIs.html>

NEW QUESTION 116

Which of the following is not a component of Elastic Beanstalk?

- A. Application
- B. Environment
- C. Docker
- D. ApplicationVersion

Answer: C

Explanation:

Answer - C
 The following are the components of Clastic Beanstalk
 1) Application - An Clastic Beanstalk application is a logical collection of Clastic Beanstalk components, including environments, versions, and environment configurations. In Clastic Beanstalk an application is conceptually similar to a folder
 2) Application version - In Clastic Beanstalk, an application version refers to a specific, labeled iteration of deployable code for a web application
 3) environment - An environment is a version that is deployed onto AWS resources. Each environment runs only a single application version at a time, however you can run the same version or different versions in many environments at the same time.
 4) environment Configuration - An environment configuration identifies a collection of parameters and settings that define how an environment and its associated resources behave.
 5) Configuration Template - A configuration template is a starting point for creating unique environment configurations. For more information on the components of Clastic beanstalk please refer to the below link <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/concepts.components.html>

NEW QUESTION 117

Which of the following is not a supported platform on Elastic Beanstalk?

- A. PackerBuilder
- B. Go
- C. Nodejs
- D. JavaSE
- E. Kubernetes

Answer: E

Explanation:

Answer-C

Below is the list of supported platforms

- *Packer Builder
- *Single Container Docker
- *Multicontainer Docker
- *Preconfigured Docker
- *Go
- *Java SE
- *Java with Tomcat
- *NET on Windows Server with IIS
- *Nodejs
- *PHP
- *Python
- *Ruby

For more information on the supported platforms please refer to the below link

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

NEW QUESTION 118

When one creates an encrypted EBS volume and attach it to a supported instance type, which of the following data types are encrypted?

Choose 3 answers from the options below

- A. Data at rest inside the volume
- B. All data copied from the EBS volume to S3
- C. All data moving between the volume and the instance
- D. All snapshots created from the volume

Answer: ACD

Explanation:

This is clearly given in the AWS documentation. Amazon EBS Encryption

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

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

For more information on EBS encryption, please refer to the below URL <http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/CBSEncryption.html>

NEW QUESTION 120

How can you resolve a dependency Error when using CloudFormation?

- A. Use the mappings attribute
- B. Use the parameter attribute
- C. Use the DependsOn attribute
- D. Use the Error attribute

Answer: C

Explanation:

The AWS troubleshooting guide for CloudFormation states the following

To resolve a dependency error, add a DependsOn attribute to resources that depend on other resources in your template. In some cases, you must explicitly declare dependencies so that AWS CloudFormation can create or delete resources in the correct order. For example, if you create an Elastic IP and a VPC with an Internet gateway in the same stack, the Elastic IP must depend on the Internet gateway attachment.

For more information on CloudFormation troubleshooting, please refer to the below URL

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

NEW QUESTION 125

Which of the following is false when it comes to using the Elastic Load balancer with OpsWorks stacks?

- A. You can attach only one load balancer to a layer.
- B. A Classic Load Balancer can span across AWS OpsWorks Stacks layers.
- C. Each load balancer can handle only one layer.
- D. You need to create the load balancer before hand and then attach it to the OpsWorks stack.

Answer: B

Explanation:

The AWS Documentation mentions the following

To use Elastic Load Balancing with a stack, you must first create one or more load balancers in the same region by using the Elastic Load Balancing console, CLI,

or API. You should be aware of the following:

You can attach only one load balancer to a layer. Each load balancer can handle only one layer.

AWS OpsWorks Stacks does not support Application Load Balancer. You can only use Classic Load Balancer with AWS OpsWorks Stacks. For more information on Classic Load Balancer with Opswork,

please visit the below url <http://docs.aws.amazon.com/opsworks/latest/userguide/layers-elb.html>

NEW QUESTION 128

You work for a startup that has developed a new photo-sharing application for mobile devices. Over recent months your application has increased in popularity; this has resulted in a decrease in the performance of the application due to the increased load. Your application has a two-tier architecture that is composed of an Auto Scaling PHP application tier and a MySQL RDS instance initially deployed with AWS CloudFormation. Your Auto Scaling group has a min value of 4 and a max value of 8. The desired capacity is now at 8 due to the high CPU utilization of the instances. After some analysis, you are confident that the performance issues stem from a constraint in CPU capacity, while memory utilization remains low. You therefore decide to move from the general-purpose M3 instances to the compute-optimized C3 instances. How would you deploy this change while minimizing any interruption to your end users?

- A. Signin to the AWS Management Console, copy the old launch configuration, and createa new launch configuration that specifies the C3 instance
- B. Update the AutoScaling group with the new launch configuratio
- C. Auto Scaling will then updatethe instance type of all running instances
- D. Signinto the AWS Management Console and update the existing launch configurationwith the newC3 instance typ
- E. Add an UpdatePolicy attribute to your AutoScaling group that specifies an AutoScaling RollingUpdate.
- F. Updatethe launch configuration specified in the AWS CloudFormation template with thenew C3 instance typ
- G. Run a stack update with the new templat
- H. Auto Scalingwill then update the instances with the new instance type.
- I. Updatethe launch configuration specified in the AWS CloudFormation template with thenew C3 instance typ
- J. Also add an UpdatePolicy attribute to your Auto Scalinggroup that specifies an AutoScalingRollingUpdat
- K. Run a stack update with thenew template

Answer: D

Explanation:

The AWS Documentation mentions the below

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 Rolling Updates for Autoscaling please see the below link:

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

NEW QUESTION 131

You are managing the development of an application that uses DynamoDB to store JSON data. You have already set the Read and Write capacity of the DynamoDB table. You are unsure of the amount of the traffic that will be received by the application during the deployment time. How can you ensure that the DynamoDB is not highly throttled and does not become a bottleneck for the application? Choose 2 answers from the options below.

- A. Monitorthe ConsumedReadCapacityUnits and ConsumedWriteCapacityUnits metric usingCloudwatch.
- B. Monitorthe SystemErrors metric using Cloudwatch
- C. Createa Cloudwatch alarm which would then send a trigger to AWS Lambda to increasethe Read and Write capacity of the DynamoDB table.
- D. Createa Cloudwatch alarm which would then send a trigger to AWS Lambda to create anew DynamoDB table.

Answer: AC

Explanation:

Refer to the following AWS Documentation that specifies what should be monitored for a DynamoDB table.

How can I determine how much of my provisioned throughput is being used?	You can monitor ConsumedReadCapacityUnits or ConsumedWriteCapacityUnits over the specified time period, to track how much of your provisioned throughput is being used.
--	---

For more information on monitoring DynamoDB please see the below link:

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

NEW QUESTION 135

When deploying applications to Elastic Beanstalk, which of the following statements is false with regards to application deployment

- A. Theapplication can be bundled in a zip file
- B. Caninclude parent directories
- C. Shouldnot exceed 512 MB in size
- D. Canbe a war file which can be deployed to the application server

Answer: B

Explanation:

The AWS Documentation mentions

When you use the AWS Clastic Beanstalk console to deploy a new application or an application version, you'll need to upload a source bundle. Your source bundle must meet the following requirements:

Consist of a single ZIP file or WAR file (you can include multiple WAR files inside your ZIP file) Not exceed 512 MB

Not include a parent folder or top-level directory (subdirectories are fine)

For more information on deploying applications to Clastic Beanstalk please see the below link: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/applications-sourcebundle.html>

NEW QUESTION 138

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

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

Answer: A

Explanation:

The AWS Documentation mentions the following

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

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

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

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

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

NEW QUESTION 140

Which of the following is not a supported platform for the Elastic beanstalk service

- A. Java
- B. AngularJS
- C. PHP
- D. .Net

Answer: B

Explanation:

Below are the supported platforms for Elastic beanstalk

Platforms

- Packer Builder
- Single Container Docker
- Multicontainer Docker
- Preconfigured Docker
- Go
- Java SE
- Java with Tomcat
- .NET on Windows Server with IIS
- Node.js
- PHP
- Python
- Ruby

For more information on Elastic beanstalk, please visit the below URL:

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

NEW QUESTION 142

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

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

Answer: C

Explanation:

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

Deployment Methods						
Method	Impact of Failed Deployment	Deploy Time	Zero Downtime	No DNS Change	Rollback Process	Code Deployed To
All at once	Downtime	⊖	×	✓	Manual Redeploy	Existing instances
Rolling	Single batch out of service; any successful batches prior to failure running new application version	⊖ ⊖ †	✓	✓	Manual Redeploy	Existing instances
Rolling with additional batch	Minimal if first batch fails, otherwise, similar to Rolling	⊖ ⊖ ⊖ †	✓	✓	Manual Redeploy	New and existing instances
Immutable	Minimal	⊖ ⊖	✓	✓	Terminate New Instances	New instances
Blue/green	Minimal	⊖ ⊖	✓	×	Swap URL	New instances

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

NEW QUESTION 145

You have a web application running on six Amazon EC2 instances, consuming about 45% of resources on each instance. You are using auto-scaling to make sure that six instances are running at all times. The number of requests this application processes is consistent and does not experience spikes. The application is critical to your business and you want high availability at all times. You want the load to be distributed evenly between all instances. You also want to use the same Amazon Machine Image (AMI) for all instances. Which of the following architectural choices should you make?

- A. Deploy 6 EC2 instances in one availability zone and use Amazon Elastic Load Balancer.
- B. Deploy 3 EC2 instances in one region and 3 in another region and use Amazon Elastic Load Balancer.
- C. Deploy 3 EC2 instances in one availability zone and 3 in another availability zone and use Amazon Elastic Load Balancer.
- D. Deploy 2 EC2 instances in three regions and use Amazon Elastic Load Balancer.

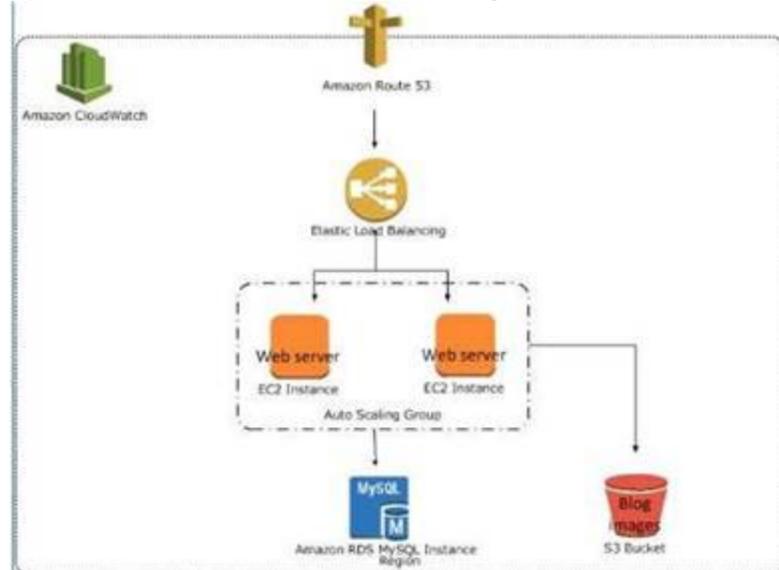
Answer: C

Explanation:

Option A is automatically incorrect because remember that the question asks for high availability. For option A, if the AZ goes down then the entire application fails. For Option B and D, the CLB is designed to only run in one region in AWS and not across multiple regions. So these options are wrong. The right option is C.

The below example shows an Elastic Loadbalancer connected to 2 EC2 instances connected via Auto Scaling. This is an example of an elastic and scalable web tier.

By scalable we mean that the Auto scaling process will increase or decrease the number of EC2 instances as required.



For more information on best practices for AWS Cloud applications, please visit the below URL:

- https://d03wsstatic.com/whitepapers/AWS_Cloud_Best_Practices.pdf

NEW QUESTION 146

You have a web application composed of an Auto Scaling group of web servers behind a load balancer, and create a new AMI for each application version for deployment. You have a new version to release, and you want to use the A/B deployment technique to migrate users over in a controlled manner while the size of the fleet remains constant over a period of 12 hours, to ensure that the new version is performing well.

What option should you choose to enable this technique while being able to roll back easily?

- A. Create an Auto scaling launch configuration with the new AMI
- B. Configure the AutoScalinggroup with the new launch configuration
- C. Use the Auto Scaling rollingupdates feature to migrate to the new version.
- D. Create an Auto Scaling launch configuration with the new AMI
- E. Create an Auto Scalinggroup configured to use the new launch configuration and to register instances with the same load balancer
- F. Vary the desired capacity of each group to migrate.
- G. Create an Auto scaling launch configuration with the new AMI
- H. Configure Auto Scaling to vary the proportion of instances launched from the two launch configurations.
- I. Create a load balancer
- J. Create an Auto Scaling launch configuration with the new AMI to use the new launch configuration and to register instances with the new load balancer
- K. Use Amazon Route 53 weighted Round Robin to vary the proportion of requests sent to the load balancers.
- L. Launch new instances using the new AMI and attach them to the Auto Scalinggroup. Configure Elastic Load Balancing to vary the proportion of requests sent to instances running the two application versions.

Answer: D

Explanation:

Since you want to control the usage to the new application in a controlled manner, the best way is to use Route53 weighted method. The AWS documentation mentions the following on this method

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 Weighted Round Robin method, please visit the link: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html#routing-policy-weighted>

NEW QUESTION 151

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

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

Answer: AEF

Explanation:

The AWS documentation mentions the following

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

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

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

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

NEW QUESTION 156

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

- A. Elastic Beanstalk
- B. Elastic Container Service
- C. OpsWorks
- D. CloudWatch

Answer: B

Explanation:

The AWS documentation provides the following information

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

For more information on ECS, please visit the link:

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

NEW QUESTION 157

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

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

Answer: BC

Explanation:

The AWS documentation mentions the below point

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

EC2 security group before you can delete the bucket or security group

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

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

NEW QUESTION 160

Which of the following are ways to secure data at rest and in transit in AWS. Choose 3 answers from the options given below

- A. Encrypt all EBS volumes attached to EC2 instances
- B. Use server-side encryption for S3
- C. Use SSL/HTTPS when using the Elastic Load Balancer
- D. Use IOPS volumes when working with EBS volumes on EC2 instances

Answer: ABC

Explanation:

The AWS documentation mentions the following

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

Data at rest inside the volume

All data moving between the volume and the instance

All snapshots created from the volume Data protection refers to protecting data while in-transit (as it travels to and from Amazon S3) and at rest (while it is stored on disks in Amazon S3 data centers). You can protect data in transit by using SSL or by using client-side encryption. You have the following options of protecting data at rest in Amazon S3.

Use Server-Side encryption - You request Amazon S3 to encrypt your object before saving it on disks in its data centers and decrypt it when you download the objects.

Use Client-Side Encryption - You can encrypt data client-side and upload the encrypted data to Amazon S3. In this case, you manage the encryption process, the encryption keys, and related tools. You can create a load balancer that uses the SSL/TLS protocol for encrypted connections (also known as SSL offload). This feature enables traffic encryption between your load balancer and the clients that initiate HTTPS sessions, and for connections between your load balancer and your L-C2 instances. For more information on securing data at rest, please refer to the below link:

- <https://dO3wsstatic.com/whitepapers/aws-securing-data-at-rest-with-encryption.pdf>

NEW QUESTION 165

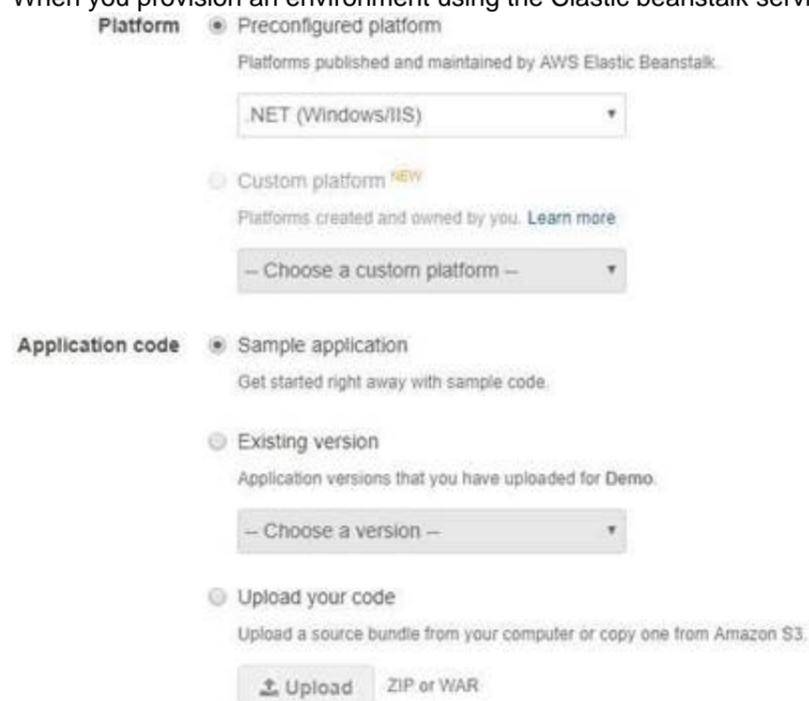
Your development team use .Net to code their web application. They want to deploy it to AWS for the purpose of continuous integration and deployment. The application code is hosted in a Git repository. Which of the following combination of steps can be used to fulfil this requirement. Choose 2 answers from the options given below

- A. Use the Elastic beanstalk service to provision an IIS platform web environment to host the application.
- B. Use the Code Pipeline service to provision an IIS environment to host the application.
- C. Create a source bundle for the .Net code and upload it as an application revision.
- D. Use a chef recipe to deploy the code and attach it to the Elastic beanstalk environment.

Answer: AC

Explanation:

When you provision an environment using the Clastic beanstalk service, you can choose the IIS platform to host the .Net based application as shown below.



You can also upload the application as a zip file and specify it as an application revision.

For more information on Elastic beanstalk and .Net environments, please refer to the below link:

http://docs^ws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_NCT.html

NEW QUESTION 166

Your company is planning on using the available services in AWS to completely automate their integration, build and deployment process. They are planning on using AWSCodeBuild to build their artefacts. When using CodeBuild, which of the following files specifies a collection of build commands that can be used by the service during the build process.

- A. appspec.yml
- B. buildspec.yml
- C. buildspec.xml
- D. appspec.json

Answer: B

Explanation:

The AWS documentation mentions the following

AWS CodeBuild currently supports building from the following source code repository providers. The source code must contain a build specification (build spec) file,

or the build spec must be declared as part of a build project definition. A buildspec\s a collection of build commands and related settings, in YAML format, that AWS

CodeBuild uses to run a build.

For more information on AWS CodeBuild, please refer to the below link: <http://docs.aws.amazon.com/codebuild/latest/userguide/planning.html>

NEW QUESTION 170

Which of the following resource is used in Cloudformation to create nested stacks

- A. AWS::CloudFormation::Stack
- B. AWS::CloudFormation::Nested
- C. AWS::CloudFormation::NestedStack
- D. AWS::CloudFormation::StackNest

Answer: A

Explanation:

The AWS Documentation mentions the following

A nested stack is a stack that you create within another stack by using the AWS:: Cloud Formation::

Stack resource. With nested stacks, you deploy and manage all resources from a single stack. You can use outputs from one stack in the nested stack group as inputs to another stack in the group

For more information on AWS::CloudFormation::Stack resource, please refer to the below link:

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

NEW QUESTION 172

Which of the following services from AWS can be integrated with the Jenkins continuous integration tool.

- A. AmazonEC2
- B. AmazonECS
- C. AmazonElastic beanstalk
- D. All of the above

Answer: D

Explanation:

The following AWS services can be integrated with Jenkins

- Amazon EC2
- Amazon ECR
- Amazon Simple Notification Service (SNS)
- Amazon ECS
- Amazon S3
- AWS CloudFormation
- AWS CodeDeploy
- AWS CodePipeline
- AWS CodeCommit
- AWS Device Farm
- AWS Elastic Beanstalk

For more information on Jenkins in AWS, please refer to the below link:

https://dOawsstatic.com/whitepapers/DevOps/Jenkins_on_AWS.pdf

NEW QUESTION 173

Your company has an application hosted on an Elastic beanstalk environment. You have been instructed that whenever application changes occur and new versions need to be deployed that the fastest deployment approach is employed. Which of the following deployment mechanisms will fulfil this requirement?

- A. All at once
- B. Rolling
- C. Immutable
- D. Rolling with batch

Answer: A

Explanation:

The following table from the AWS documentation shows the deployment time for each deployment methods.

Method	Impact of Failed Deployment	Deploy Time	Zero Downtime	No DNS Change	Rollback Process	Code Deployed To
All at once	Downtime	☹	X	✓	Re-deploy	Existing instances
Rolling	Single batch out of service. Any successful batches prior to failure running new application version.	☹ ☹ ↑	✓	✓	Re-deploy	Existing instances
Rolling with additional batch	Minimal if first batch fails, otherwise similar to Rolling.	☹ ☹ ☹ ↑	✓	✓	Re-deploy	New & existing instances
Immutable	Minimal	☹ ☹ ☹ ☹	✓	✓	Re-deploy	New instances
Blue/green	Minimal	☹ ☹ ☹ ☹	✓	X	Swap URL	New instances

For more information on Elastic beanstalk deployments, please refer to the below link: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.deploy-existing-version.html>

NEW QUESTION 178

In AWS Code Deploy which of the following deployment types are available. Choose 2 answers from the options given below

- A. In-placedeployments
- B. Rollingdeployments
- C. Immutabledeployments
- D. Blue/Greendeployments

Answer: AD

Explanation:

The AWS documentation mentions the following

Deployment type: The method used to make the latest application revision available on instances in a deployment group.

In-place deployment: The application on each instance in the deployment group is stopped, the latest application revision is installed, and the new version of the application is started and validated. You can choose to use a load balancer so each instance is deregistered during its deployment and then restored to service after the deployment is complete.

Blue/green deployment: The instances in a deployment group (the original environment) are replaced by a different set of instances (the replacement environment) using these steps:

Instances are provisioned for the replacement environment.

o The latest application revision is installed on the replacement instances,

o An optional wait time occurs for activities such as application testing and system verification. Instances in the replacement environment are registered with an Elastic Load Balancing load balancer, causing traffic to be rerouted to them. Instances in the original environment are deregistered and can be terminated or kept running for other uses. For more information on the components of AWS Code Deploy, please refer to the below link:

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

NEW QUESTION 182

You are using Elastic beanstalk to deploy an application that consists of a web and application server. There is a requirement to run some python scripts before the application version is deployed to the web server. Which of the following can be used to achieve this?

- A. Makeuse of container commands
- B. Makeuse of Docker containers
- C. Makeuse of custom resources
- D. Makeuse of multiple elastic beanstalk environments

Answer: A

Explanation:

The AWS Documentation mentions the following

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. For more information on Container commands, please visit the below URL: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/customize-containers-ec2.html>

NEW QUESTION 185

Your company has a set of resources hosted in AWS. Your IT Supervisor is concerned with the costs being incurred with the current set of AWS resources and wants to monitor the cost usage. Which of the following mechanisms can be used to monitor the costs of the AWS resources and also look at the possibility of cost optimization. Choose 3 answers from the options given below

- A. Use the Cost Explorer to see the costs of AWS resources
- B. Create budgets in billing section so that budgets are set beforehand
- C. Send all logs to Cloudwatch logs and inspect the logs for billing details
- D. Consider using the Trusted Advisor

Answer: ABD

Explanation:

The AWS Documentation mentions the following

1) For a quick, high-level analysis use Cost Explorer, which is a free tool that you can use to view graphs of your AWS spend data. It includes a variety of filters and preconfigured views, as well as forecasting capabilities. Cost Explorer displays data from the last 13 months, the current month, and the forecasted costs for the next three months, and it updates this data daily.

2) Consider using budgets if you have a defined spending plan for a project or service and you want to track how close your usage and costs are to exceeding your budgeted amount. Budgets use data from Cost Explorer to provide you with a quick way to see your usage-to-date and current estimated charges from AWS. You can also set up notifications that warn you if you exceed or are about to exceed your budgeted amount.

3) Visit the AWS Trusted Advisor console regularly. Trusted Advisor works like a customized cloud expert, analyzing your AWS environment and providing best practice recommendations to help you save money, improve system performance and reliability, and close security gaps.

For more information on cost optimization, please visit the below URL:

- <https://aws.amazon.com/answers/account-management/cost-optimization-monitor/>

NEW QUESTION 188

Your development team is developing a mobile application that access resources in AWS. The users accessing this application will be logging in via Facebook and Google. Which of the following AWS mechanisms would you use to authenticate users for the application that needs to access AWS resources

- A. Use separate IAM users that correspond to each Facebook and Google user
- B. Use separate IAM Roles that correspond to each Facebook and Google user
- C. Use Web identity federation to authenticate the users
- D. Use AWS Policies to authenticate the users

Answer: C

Explanation:

The AWS documentation mentions the following

You can directly configure individual identity providers to access AWS resources using web identity federation. AWS currently supports authenticating users using web identity federation through several identity providers: Login with Amazon

Facebook Login

Google Sign-in For more information on Web identity federation please visit the below URL:

- <http://docs.aws.amazon.com/sdk-for-javascript/v2/developer-guide/load-ing-browser-credentials-federated-id.htm>

NEW QUESTION 189

In reviewing the Auto-Scaling events for your application you notice that your application is scaling up and down multiple times in the same hour. What design choice could you make to optimize for costs while preserving elasticity?

Choose 2 options from the choices given below

- A. Modify the Auto Scaling policy to use scheduled scaling actions
- B. Modify the Auto Scaling Group cool down timers
- C. Modify the Amazon Cloudwatch alarm period that triggers your Auto Scaling scale down policy.
- D. Modify the Auto Scaling group termination policy to terminate the newest instance first.

Answer: BC

Explanation:

The Auto Scaling cooldown period is a configurable setting for your Auto Scaling group that helps to ensure that Auto Scaling doesn't launch or terminate additional instances before the previous scaling activity takes effect. After the Auto Scaling group dynamically scales using a simple scaling policy, Auto Scaling waits for the cooldown period to complete before resuming scaling activities. When you manually scale your Auto Scaling group, the default is not to wait for the cooldown period,

but you can override the default and honor the cooldown period. Note that if an instance becomes

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

For more information on Autoscale cool down timers please visit the URL: <http://docs.aws.amazon.com/autoscaling/latest/userguide/Cooldown.htm>

You can also modify the Cloudwatch triggers to ensure the thresholds are appropriate for the scale down policy For more information on Autoscaling user guide please visit the URL: <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scale-based-on-demand.html>

NEW QUESTION 190

There is a company website that is going to be launched in the coming weeks. There is a probability that the traffic will be quite high in the first couple of weeks. In the event of a load failure, how can you set up DNS failover to a static website? Choose the correct answer from the options given below.

- A. Duplicate the exact application architecture in another region and configure DNS weight-based routing
- B. Enable failover to an on-premise data center to the application hosted there.
- C. Use Route 53 with the failover option to failover to a static S3 website bucket or CloudFront distribution.
- D. Add more servers in case the application fails.

Answer: C

Explanation:

Amazon Route 53 health checks monitor the health and performance of your web applications, web servers, and other resources.

If you have multiple resources that perform the same function, you can configure DNS failover so that Amazon Route 53 will route your traffic from an unhealthy resource to a healthy resource. For example, if you have two web servers and one web server becomes unhealthy, Amazon Route 53 can route traffic to the other web server. So you can route traffic to a website hosted on S3 or to a CloudFront distribution.

For more information on DNS failover using Route 53, please refer to the below link:

- <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.htm>

NEW QUESTION 193

A company is running three production web server reserved EC2 instances with EBS-backed root volumes. These instances have a consistent CPU load of 80%. Traffic is being distributed to these instances by an Elastic Load Balancer. They also have production and development Multi-AZ RDS MySQL databases. What recommendation would you make to reduce cost in this environment without affecting availability of mission-critical systems? Choose the correct answer from the options given below

- A. Consider using on-demand instances instead of reserved EC2 instances
- B. Consider not using a Multi-AZ RDS deployment for the development database
- C. Consider using spot instances instead of reserved EC2 instances
- D. Consider removing the Elastic Load Balancer

Answer: B

Explanation:

Multi-AZ databases is better for production environments rather than for development environments, so you can reduce costs by not using this for development environments. 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. Since the endpoint for your DB Instance remains the same after a failover, your application can resume database operation without the need for manual administrative intervention

For more information on Multi-AZ RDS, please refer to the below link: <https://aws.amazon.com/rds/details/multi-az/>

NEW QUESTION 196

Your company is concerned with EBS volume backup on Amazon EC2 and wants to ensure they have proper backups and that the data is durable. What solution would you implement and why? Choose the correct answer from the options below

- A. Configure Amazon Storage Gateway with EBS volumes as the data source and store the backups on-premise through the storage gateway

- B. Write a cronjob on the server that compresses the data that needs to be backed up using gzip compression, then use AWS CLI to copy the data into an S3 bucket for durability
- C. Use a lifecycle policy to back up EBS volumes stored on Amazon S3 for durability
- D. Write a cronjob that uses the AWS CLI to take a snapshot of production EBS volume
- E. The data is durable because EBS snapshots are stored on the Amazon S3 standard storage class

Answer: D

Explanation:

You can take snapshots of EBS volumes and to automate the process you can use the CLI. The snapshots are automatically stored on S3 for durability. For more information on EBS snapshots, please refer to the below link: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSSnapshots.html>

NEW QUESTION 201

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

DOP-C01 Practice Exam Features:

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