DOP-C01 Dumps

AWS Certified DevOps Engineer- Professional

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

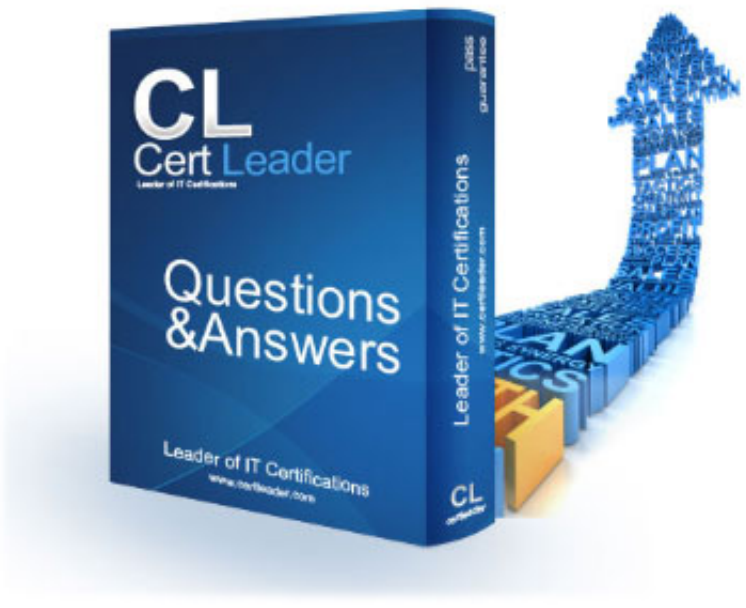

NEW QUESTION 1
You have an application which consists of EC2 instances in an Auto Scaling group. Between a particular time frame every day, there is an increase in traffic to your website. Hence users are complaining of a poor response time on the application. You have configured your Auto Scaling group to deploy one new EC2 instance when CPU utilization is greater than 60% for 2 consecutive periods of 5 minutes. What is the least cost-effective way to resolve this problem?

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

Answer: B

Explanation:
If you increase the minimum number of instances, then they will be running even though the load is not high on the website. Hence you are incurring cost even though there is no need.
All of the remaining options are possible options which can be used to increase the number of instances on a high load.
For more information on On-demand scaling, please refer to the below link: http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scale-based-on-demand.html
Note: The tricky part where the question is asking for 'least cost effective way". You got the design consideration correctly but need to be careful on how the question is phrased.

NEW QUESTION 2
You have an application running a specific process that is critical to the application's functionality, and have added the health check process to your Auto Scaling Group. The instances are showing healthy but the application itself is not working as it should. What could be the issue with the health check, since it is still showing the instances as healthy.

A. You do not have the time range in the health check properly configured
B. It is not possible for a health check to monitor a process that involves the application
C. The health check is not configured properly
D. The health check is not checking the application process

Answer: D

Explanation:
If you have custom health checks, you can send the information from your health checks to Auto Scaling so that Auto Scaling can use this information. For example, if
you determine that an instance is not functioning as expected, you can set the health status of the instance to Unhealthy. The next time that Auto Scaling performs a
health check on the instance, it will determine that the instance is unhealthy and then launch a replacement instance
For more information on Autoscaling health checks, please refer to the below document link: from AWS
http://docs.aws.amazon.com/autoscaling/latest/userguide/healthcheck.html

NEW QUESTION 3
You have deployed an application to AWS which makes use of Autoscaling to launch new instances. You now want to change the instance type for the new instances. Which of the following is one of the action items to achieve this deployment?

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

Answer: C

Explanation:
The ideal way is to create a new launch configuration, attach it to the existing Auto Scaling group, and terminate the running instances.
Option A is invalid because Clastic beanstalk cannot launch new instances on demand. Since the current scenario requires Autoscaling, this is not the ideal option
Option B is invalid because this will be a maintenance overhead, since you just have an Autoscaling Group. There is no need to create a whole Cloudformation template for this.
Option D is invalid because Autoscaling Group will still launch CC2 instances with the older launch configuration
For more information on Autoscaling Launch configuration, please refer to the below document link: from AWS
http://docs.aws.amazon.com/autoscaling/latest/userguide/l_aunchConfiguration.html

NEW QUESTION 4
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 Ad dTo Load Balancer, Auto Scaling launches the instances but does not add them to the load balancer or target group. If you resume
the AddTo Load Balancer process. Auto Scaling resumes adding instances to the load balancer or target group when they are launched. However, Auto Scaling does
not add the instances that were launched while this process was suspended. You must register those
instances manually.

Option A is invalid because this just balances the number of CC2 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 CC2 or Clastic 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 5
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 queu
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 triggere
C. Create an Amazon EC2 Auto Scaling group with both minimum and desired Instances configured to 0. Worker nodes in thisgroup spawn when messages are added to the queu
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 topi
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 triggere
J. Create an HTTP subscriber to this topic that notifies your application via HTTP POST when an alarm is triggere
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 triggere
N. Create an HTTP subscriber to this topic that notifies your application via HTTP POST when an alarm is triggere
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 fulfils all the requirements
1) First is to create a SNS topic for each group so that the required members get the email addresses.
2) Ensure the application uses the HTTPS endpoint and the SDK to publish messages Option A is invalid because the SQS service is not required.
Option B and C are incorrect. As per the requirement we need to provide notification to only those on-call teams who are working in that particular shift when an alarm is triggered. It need not have to be send to all the on-call teams of the company. With Option B & C, since we are not configuring the SNS topic for each on call team the notifications will be send to all the on-call teams. Hence these 2 options are invalid. For more information on setting up notifications, please refer to the below document link: from AWS http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US_SetupSNS.html

NEW QUESTION 6
You are responsible for your company's large multi-tiered Windows-based web application running on Amazon EC2 instances situated behind a load balancer. While reviewing metrics, you've started noticing an upwards trend for slow customer page load time. Your manager has asked you to come up with a solution to ensure that customer load time is not affected by too many requests per second. Which technique would you use to solve this issue?

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

Answer: C

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

NEW QUESTION 7
Management has reported an increase in the monthly bill from Amazon Web Services, and they are extremely concerned with this increased cost. Management has asked you to determine the exact cause of this increase. After reviewing the billing report, you notice an increase in the data transfer cost. How can you provide management with a better insight into data transfer use?

A. Update your Amazon CloudWatch metrics to use five-second granularity, which will give better detailed metrics that can be combined with your billing data to pinpoint anomalies.
B. Use Amazon CloudWatch Logs to run a map-reduce on your logs to determine high usage and data transfer.
C. Deliver custom metrics to Amazon CloudWatch per application that breaks down application data transfer into multiple, more specific data points.D- Using

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

Answer: C

Explanation:
You can publish your own metrics to CloudWatch using the AWS CLI or an API. You can view statistical graphs of your published metrics with the AWS Management Console.
CloudWatch stores data about a metric as a series of data points. Each data point has an associated time stamp. You can even publish an aggregated set of data points called a statistic set.
If you have custom metrics specific to your application, you can give a breakdown to the management on the exact issue.
Option A won't be sufficient to provide better insights.
Option B is an overhead when you can make the application publish custom metrics Option D is invalid because just the ELB metrics will not give the entire picture
For more information on custom metrics, please refer to the below document link: from AWS
http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.html

NEW QUESTION 8
You have an application consisting of a stateless web server tier running on Amazon EC2 instances behind load balancer, and are using Amazon RDS with read replicas. Which of the following methods should you use to implement a self-healing and cost-effective architecture? Choose 2 answers from the optionsgiven 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 Scalinggroup 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 Scalinggroup 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 Scalinggroup 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 CC2 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). Cach 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 9
You are using Elastic Beanstalk to manage your e-commerce store. The store is based on an open source e- commerce platform and is deployed across multiple instances in an Auto Scaling group. Your development team often creates new "extensions" for the e-commerce store. These extensions include PHP source code as well as an SQL upgrade script used to make any necessary updates to the database schema. You have noticed that some extension deployments fail due to an error when running the SQL upgrade script. After further investigation, you realize that this is because the SQL script is being executed on all of your Amazon EC2 instances. How would you ensure that the SQL script is only executed once per deployment regardless of how many Amazon EC2 instances are running at the time?

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

Answer: A

Explanation:
You can use the container_commands key to execute commands that affect your application source code. Container commands run after the application and web server have been set up and the application version archive has been extracted, but before the application version is deployed. Non-container commands and other customization operations are performed prior to the application source code being extracted.
You can use leader_only to only run the command on a single instance, or configure a test to only run the command when a test command evaluates to true. Leader-only container commands are only executed during environment creation and deployments, while other commands and server customization operations are performed every time an instance is provisioned or updated. Leader- only container commands are not executed due to launch configuration changes, such as a change in the AMI Id or instance type. For more information on customizing containers, please visit the below URL:
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/customize-containers-ec2.html

NEW QUESTION 10
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 internalAmazon 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 10
As part of your continuous deployment process, your application undergoes an I/O load performance test before it is deployed to production using new AMIs. The application uses one Amazon Elastic Block Store (EBS) PIOPS volume per instance and requires consistent I/O performance. Which of the following must be carried out to ensure that I/O load performance tests yield the correct results in a repeatable manner?

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

Answer: B

Explanation:
During the AMI-creation process, Amazon CC2 creates snapshots of your instance's root volume and any other CBS volumes attached to your instance
New CBS volumes receive their maximum performance the moment that they are available and do not require initialization (formerly known as pre-warming). However, storage blocks on volumes that were restored from snapshots must to initialized (pulled
down from Amazon S3 and written to the volume) before you can access the block. This preliminary action takes time and can cause a significant increase in the latency of an I/O operation the first time each block is accessed. For most applications, amortizing this cost over the lifetime of the volume is acceptable.
Option A is invalid because block sizes are predetermined and should not be randomly selected. Option C is invalid because this is part of continuous integration and hence volumes can be destroyed after the test and hence there should not be snapshots created unnecessarily
Option D is invalid because the encryption is a security feature and not part of load tests normally. For more information on CBS initialization please refer to the below link:
• http://docs.aws.a mazon.com/AWSCC2/latest/UserGuide/ebs-in itialize.html

NEW QUESTION 12
You have a large number of web servers in an Auto Scalinggroup behind a load balancer. On an hourly basis, you want to filter and process the logs to collect data on unique visitors, and then put that data in a durable data store in order to run reports. Web servers in the Auto Scalinggroup are constantly launching and terminating based on your scaling policies, but you do not want to lose any of the log data from these servers during a stop/termination initiated by a user or by Auto Scaling. What two approaches will meet these requirements? Choose two answers from the optionsgiven below.

A. Install an Amazon Cloudwatch Logs Agent on every web server during the bootstrap proces
B. Create a CloudWatch log group and defineMetric Filters to create custom metrics that track unique visitors from the streaming web server log
C. Create a scheduled task on an Amazon EC2 instance that runs every hour to generate a new report based on the Cloudwatch custom metric
D. ^/
E. On the web servers, create a scheduled task that executes a script to rotate and transmit the logs to Amazon Glacie
F. Ensure that the operating system shutdown procedure triggers a logs transmission when the Amazon EC2 instance is stopped/terminate
G. Use Amazon Data Pipeline to process the data in Amazon Glacier and run reports every hour.
H. On the web servers, create a scheduled task that executes a script to rotate and transmit the logs to an Amazon S3 bucke
I. Ensure that the operating system shutdown procedure triggers a logs transmission when the Amazon EC2 instance is stopped/terminate
J. Use AWS Data Pipeline to move log data from the Amazon S3 bucket to Amazon Redshift In order to process and run reports every hour.
K. Install an AWS Data Pipeline Logs Agent on every web server during the bootstrap proces
L. Create a log group object in AWS Data Pipeline, and define Metric Filters to move processed log data directly from the web servers to Amazon Redshift and run reports every hour.

Answer: AC

Explanation:
You can use the Cloud Watch Logs agent installer on an existing CC2 instance to install and configure the Cloud Watch Logs agent.
For more information, please visit the below link:
• http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/Qu ickStartCC2Instance.html
You can publish your own metrics to Cloud Watch using the AWS CLI or an API. For more information, please visit the below link:

• http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.htmI Amazon Redshift is a fast, fully managed data warehouse that makes it simple and cost-effective to analyze all your data using standard SQL and your existing Business Intelligence (BI) tools. It allows you to run complex analytic queries against petabytes of structured data, using sophisticated query optimization, columnar storage on high-performance local disks, and massively parallel query execution. Most results come back in seconds. For more information on copying data from S3 to redshift, please refer to the below link:
• http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/dp-copydata- redshift html

NEW QUESTION 13
You have a set of EC2 instances hosted in AWS. You have created a role named DemoRole and assigned that role to a policy, but you are unable to use that role with an instance. Why is this the case.

A. You need to create an instance profile and associate it with that specific role.
B. You are not able to associate an 1AM role with an instanceC You won't be able to use that role with an instance unless you also create a user and associate it with that specific role
C. You won't be able to use that role with an instance unless you also create a usergroup and associate it with that specific role.

Answer: A

Explanation:
An instance profile is a container for an 1AM role that you can use to pass role information to an CC2 instance when the instance starts.
Option B is invalid because you can associate a role with an instance
Option C and D are invalid because using users or user groups is not a pre-requisite For more information on instance profiles, please visit the link:
• http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-ro le-ec2_instance- profiles.htm I

NEW QUESTION 16
You are using Elastic Beanstalk to manage your application. You have a SQL script that needs to only be executed once per deployment no matter how many EC2 instances you have running. How can you do this?

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

Answer: C

Explanation:
You can use the container_commands key to execute commands that affect your application source code. Container commands run after the application and web server have been set up and the application version archive has been extracted, but before the application version is deployed. Non- container commands and other customization operations are performed prior to the application source code being extracted.
You can use leader_only to only run the command on a single instance, or configure a test to only run the command when a test command evaluates to true. Leader-only container commands are only executed during environment creation and deployments, while other commands and server customization operations are performed every time an instance is provisioned or updated. Leader- only container commands are not executed due to launch configuration changes, such as a change in the AMI Id or instance type. For more information on customizing containers, please visit the below URL:
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/customize-containers-ec2.html

NEW QUESTION 18
You work for an insurance company and are responsible for the day-to-day operations of your company's online quote system used to provide insurance quotes to members of the public. Your company wants to use the application logs generated by the system to better understand customer behavior. Industry, regulations also require that you retain all application logs for the system indefinitely in order to investigate fraudulent claims in the future. You have been tasked with designing a log management system with the following requirements:
- All log entries must be retained by the system, even during unplanned instance failure.
- The customer insight team requires immediate access to the logs from the past seven days.
- The fraud investigation team requires access to all historic logs, but will wait up to 24 hours before these logs are available.
How would you meet these requirements in a cost-effective manner? Choose three answers from the options below

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

Answer: CEF

Explanation:
Since all logs need to be stored indefinitely. Glacier is the best option for this. One can use Lifecycle events to stream the data from S3 to Glacier
Lifecycle configuration enables you to specify the lifecycle management of objects in a bucket. The configuration is a set of one or more rules, where each rule defines an action for Amazon S3 to apply to a group of objects. These actions can be classified as
follows:
• Transition actions - In which you define when objects transition to another storage class. For example, you may choose to transition objects to the STANDARDJA QK for infrequent access) storage class 30 days after creation, or archive objects to the GLACIER storage class one year after creation.
• Expiration actions - In which you specify when the objects expire. Then Amazon S3 deletes the expired objects on your behalf. For more information on Lifecycle events, please refer to the below link:
• http://docs.aws.a mazon.com/AmazonS3/latest/dev/object-lifecycle-mgmt.htm I You can use scripts to put the logs onto a new volume and then transfer those logs to S3.

Note:
Moving the logs from CBS volume to S3 we have some custom scripts running in the background. Inorder to ensure the minimum memory requirements for the OS and the applications for the script to execute we can use a cost effective ec2 instance.
Considering the computing resource requirements of the instance and the cost factor a tZmicro instance can be used in this case.
The following link provides more information on various t2 instances. https://docs.aws.amazon.com/AWSCC2/latest/WindowsGuide/t2-instances.html
Question is "How would you meet these requirements in a cost-effective manner? Choose three answers from the options below"
So here user has to choose the 3 options so that the requirement is fulfilled. So in the given 6 options, options C, C and F fulfill the requirement.
" The CC2s use CBS volumes and the logs are stored on CBS volumes those are marked for non- termination" - is one of the way to fulfill requirement. So this shouldn't be a issue.

NEW QUESTION 20
You need to implement Blue/Green Deployment for several multi-tier web applications. Each of them has Its Individual infrastructure:
Amazon Elastic Compute Cloud (EC2) front-end servers, Amazon ElastiCache clusters, Amazon Simple Queue Service (SQS) queues, and Amazon Relational Database (RDS) Instances.
Which combination of services would give you the ability to control traffic between different deployed versions of your application?

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

Answer: B

Explanation:
This an example of Blue green deployment

With Amazon Route 53, you can define a percentage of traffic to go to the green environment and gradually update the weights until the green environment carries the full production traffic. A weighted distribution provides the ability to perform canary analysis where a small percentage of production traffic is introduced to a new environment. You can test the new code and monitor for errors, limiting the blast radius if any issues are encountered. It also allows the green environment to scale out to support the full production load if you're using Elastic Load Balancing.
When if s time to promote the green environment/stack into production, update DNS records to point to the green environment/stack's load balancer. You can also do this DNS flip gradually by using the Amazon Route 53 weighted routing policy. For more information on Blue green deployment, please refer to the link:
• https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 25
You've been tasked with improving the current deployment process by making it easier to deploy and reducing the time it takes. You have been tasked with creating a continuous integration (CI) pipeline that can build AMI'S. Which of the below is the best manner to get this done. Assume that at max your development team will be deploying builds 5 times a week.

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

Answer: D

Explanation:
Since the number of calls is just a few times a week, there are many open source systems such as Jenkins which can be used as CI based systems.
Jenkins can be used as an extensible automation server, Jenkins can be used as a simple CI server or turned into the continuous delivery hub for any project.
For more information on the Jenkins CI tool please refer to the below link:
• https://jenkins.io/
Option A and C are partially correct, but since you just have 5 deployments per week, having separate instances which consume costs is not required. Option B is partially correct, but again having a separate system such as Opswork for such a low number of deployments is not required.

NEW QUESTION 29

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

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

Answer: AB

Explanation:
Amazon Cloud Watch monitors your Amazon Web Services (AWS) resources and the applications you run on AWS in real time. You can use Cloud Watch to collect and track metrics, which are variables you can measure for your resources and applications. Cloud Watch alarms send notifications or automatically make changes to the resources you are monitoring based on rules that you define.
For more information on AWS Cloudwatch, please refer to the below document link: from AWS
• http://docs.aws.a mazon.com/AmazonCloudWatch/latest/monitoring/WhatlsCloud Watch.htm I Amazon Cloud Watch uses Amazon SNS to send email. First, create and subscribe to an SNS topic.
When you create a CloudWatch alarm, you can add this SNS topic to send an email notification when the alarm changes state
For more information on AWS Cloudwatch and SNS, please refer to the below document link: from AWS
http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US_SetupSNS.html

NEW QUESTION 30
You are designing a system which needs, at a minimum, 8 m4.large instances operating to service traffic. When designing a system for high availability in the us-east-1 region, which has 6 Availability Zones, your company needs to be able to handle the death of a full availability zone. How should you distribute the servers, to save as much cost as possible, assuming all of the EC2 nodes are properly linked to an ELB? Your VPC account can utilize us-east-1's AZ's a through f, inclusive.

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

Answer: C

Explanation:
The best way is to distribute the instances across multiple AZ's to get the best and avoid a disaster scenario. With this scenario, you will always a minimum of more than 8 servers even if one AZ were to go down. Even though A and D are also valid options, the best option when it comes to distribution is Option C. For more information on High Availability and Fault tolerance, please refer to the below link:
https://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_ftha_04.pdf

NEW QUESTION 31
You have an application hosted in AWS. You wanted to ensure that when certain thresholds are reached, a Devops Engineer is notified. Choose 3 answers from the options given below

A. Use CloudWatch Logs agent to send log data from the app to CloudWatch Logs from Amazon EC2 instances
B. Pipe data from EC2 to the application logs using AWS Data Pipeline and CloudWatch
C. Once a CloudWatch alarm is triggered, use SNS to notify the Senior DevOps Engineer.
D. Set the threshold your application can tolerate in a CloudWatch Logs group and link a CloudWatch alarm on that threshold.

Answer: ACD

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

NEW QUESTION 32
Your company releases new features with high frequency while demanding high application availability. As part of the application's A/B testing, logs from each updated Amazon EC2 instance of the application need to be analyzed in near real-time, to ensure that the application is working
flawlessly after each deployment. If the logs show any anomalous behavior, then the application version of the instance is changed to a more stable one. Which of the following methods should you use for shipping and analyzing the logs in a highly available manner?

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

Answer: C

Explanation:
Answer - C
You can use Kinesis Streams for rapid and continuous data intake and aggregation. The type of data used includes IT infrastructure log data, application logs,

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

NEW QUESTION 33
You need to deploy a Node.js application and do not have any experience with AWS. Which deployment method will be the simplest for you to deploy?

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

Answer: A

Explanation:
With Elastic Beanstalk, you can quickly deploy and manage applications in the AWS Cloud without worrying about the infrastructure that runs those applications. AWS Elastic Beanstalk reduces management complexity without restricting choice or control. You simply upload your application, and Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring
For more information on Elastic beanstalk please refer to the below link:
• http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/Welcome.html

NEW QUESTION 38
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 41
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 Clastic 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 CC2 Instance
Option D is invalid because there is a limit of 5 VPCs per region. For more information on Clastic IP's, please visit the below URL:
• http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/elastic-i p-addresses-eip.html

NEW QUESTION 43
You are designing a service that aggregates clickstream data in batch and delivers reports to subscribers via email only once per week. Data is extremely spikey, geographically distributed, high- scale, and unpredictable. How should you design this system?

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

Answer: B

Explanation:
When you look at building reports or analyzing data from a large data set, you need to consider CMR because this service is built on the Hadoop framework which is used to processes large data sets.
The ideal approach to getting data onto CMR is to use S3. Since the Data is extremely spikey and geographically distributed, using edge locations via Cloudfront distributions is the best way to fetch the data.
Option A is invalid because RedShift is more of a petabyte storage cluster.
Option C is invalid because having both Kinesis and CMR for the job analysis is redundant. Option D is invalid because Elastic Search is not an option for processing records.
For more information on Amazon CMR, please visit the below URL:
• https://aws.amazon.com/emr/

NEW QUESTION 45
You are building a game high score table in DynamoDB. You will store each user's highest score for each game, with many games, all of which have relatively similar usage levels and numbers of players. You need to be able to look up the highest score for any game. What's the best DynamoDB key structure?

A. HighestScore as the hash/only key.
B. GamelD as the hash key, HighestScore as the range ke
C. GamelD as the hash/only key.
D. GamelDastherange/onlykey.

Answer: B

Explanation:
It always best to choose the hash key as the column that will have a wide range of values. This is also given in the AWS documentation
Choosing a Partition Key
The following table compares some common partition key schemas for provisioned throughput efficiency:

Next since you need to sort by the Highest Score, you need to use that as the sort key For more information on Table Guidelines, please visit the below URL:
• http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Guide linesForTables.html

NEW QUESTION 49
Your application consists of 10% writes and 90% reads. You currently service all requests through a Route53 Alias Record directed towards an AWS ELB, which sits in front of an EC2 Auto Scaling Group. Your system isgetting very expensive when there are large traffic spikes during certain news events, during which many more people request to read similar data all at the same time. What is the simplest and cheapest way to reduce costs and scale with spikes like this?

A. Create an S3 bucket and asynchronously replicate common requests responses into S3 object
B. When a request comes in for a precomputed response, redirect to AWS S3.
C. Create another ELB and Auto Scaling Group layer mounted on top of the other system, adding a tier to the syste
D. Serve most read requests out of the top layer.
E. Create a CloudFront Distribution and direct Route53 to the Distributio
F. Use the ELB as an Origin and specify Cache Behaviours to proxy cache requests which can be served late.
G. Create a Memcached cluster in AWS ElastiCach
H. Create cache logic to serve requests which can be served late from the in-memory cache for increased performance.

Answer: C

Explanation:
Use Cloudf rant distribution for distributing the heavy reads for your application. You can create a
zone apex record to point to the Cloudfront distribution.
You can control how long your objects stay in a CloudFront cache before CloudFront forwards another request to your origin. Reducing the duration allows you to serve dynamic content. Increasing the duration means your users get better performance because your objects are more likely to be served directly from the edge cache. A longer duration also reduces the load on your origin.
For more information on Cloudfront object expiration, please visit the below URL:
http://docs.aws.amazon.com/AmazonCloudFrant/latest/DeveloperGuide/Cxpiration.html

NEW QUESTION 52
You need the absolute highest possible network performance for a cluster computing application. You already selected homogeneous instance types supporting 10 gigabit enhanced networking, made sure that your workload was network bound, and put the instances in a placement group. What is the last optimization you can make?

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

Answer: A

Explanation:
Jumbo frames allow more than 1500 bytes of data by increasing the payload size per packet, and thus increasing the percentage of the packet that is not packet overhead. Fewer packets are needed to send the same amount of usable data. However, outside of a given AWS region (CC2-Classic), a single VPC, or a VPC peering
connection, you will experience a maximum path of 1500 MTU. VPN connections and traffic sent over an Internet gateway are limited to 1500 MTU. If packets are over
1500 bytes, they are fragmented, or they are dropped if the Don't Fragment flag is set in the IP header.
For more information on Jumbo Frames, please visit the below URL:
http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/network_mtu.htm#jumbo_frame_instance s

NEW QUESTION 54
Your application's Auto Scaling Group scales up too quickly, too much, and stays scaled when traffic decreases. What should you do to fix this?

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

Answer: B

Explanation:
The ideal case is that the right metric is not being used for the scale up and down.
Option A is not valid because it mentions that the cooldown is not happening when the traffic decreases, that means the metric threshold for the scale down is not occurring in Cloudwatch
Option C is not valid because increasing the Cloudwatch alarm metric will not ensure that the instances scale down when the traffic decreases.
Option D is not valid because the question does not mention any constraints that points to the instance size. For an example on using custom metrics for scaling in and out, please follow the below link for a use case.
• https://blog.powerupcloud.com/aws-autoscaling-based-on-database-query-custom-metrics- f396c16e5e6a

NEW QUESTION 55
You need to run a very large batch data processingjob one time per day. The source data exists
entirely in S3, and the output of the processingjob should also be written to S3 when finished. If you need to version control this processingjob and all setup and teardown logic for the system, what approach should you use?.

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

Answer: B

Explanation:
With AWS Cloud Formation, you can update the properties for resources in your existing stacks.
These changes can range from simple configuration changes, such
as updating the alarm threshold on a Cloud Watch alarm, to more complex changes, such as updating the Amazon Machine Image (AMI) running on an Amazon EC2
instance. Many of the AWS resources in a template can be updated, and we continue to add support for more.
For more information on Cloudformation version control, please visit the below URL:
http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/updating.stacks.wa I kthrough.htm I

NEW QUESTION 57
What is required to achieve gigabit network throughput on EC2? You already selected cluster- compute, 10GB instances with enhanced networking, and your workload is already network-bound, but you are not seeing 10 gigabit speeds.

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

Answer: D

Explanation:
A placement group is a logical grouping of instances within a single Availability Zone. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. To provide the lowest latency, and the highest packet-per-second network performance for your placement group, choose an instance type that supports enhanced networking. For more information on Placement Groups, please visit the below URL:
http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/placement-groups.html

NEW QUESTION 62
Your team wants to begin practicing continuous delivery using CloudFormation, to enable automated builds and deploys of whole, versioned stacks or stack layers.

You have a 3-tier, mission-critical system. Which of the following is NOT a best practice for using CloudFormation in a continuous delivery environment?

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

Answer: B

Explanation:
Answer - B
Some of the best practices for Cloudformation are
• Created Nested stacks
As your infrastructure grows, common patterns can emerge in which you declare the same components in each of your templates. You can separate out these common components and create dedicated templates for them. That way, you can mix and match different templates but use nested stacks to create a single, unified stack. Nested stacks are stacks that create other stacks. To create nested stacks, use the AWS::CloudFormation::Stackresource in your template to reference other templates.
• Reuse Templates
After you have your stacks and resources set up, you can reuse your templates to replicate your infrastructure in multiple environments. For example, you can create environments for development, testing, and production so that you can test changes before implementing them into production. To make templates reusable, use the parameters, mappings, and conditions sections so that you can customize your stacks when you create them. For example, for your development environments, you can specify a lower-cost instance type compared to your production environment, but all other configurations and settings remain the same. For more information on Cloudformation best practises, please visit the below URL:
http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/best-practices.html

NEW QUESTION 66
You run a 2000-engineer organization. You are about to begin using AWS at a large scale for the first time. You want to integrate with your existing identity management system running on Microsoft Active Directory, because your organization is a power-user of Active Directory. How should you manage your AWS identities in the most simple manner?

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

Answer: B

Explanation:
AD Connector is a directory gateway with which you can redirect directory requests to your on- premises Microsoft Active Directory without caching any information in the cloud. AD Connector comes in two sizes, small and large. A small AD Connector is designed for smaller organizations of up to 500 users. A large AD Connector
can support larger organizations of up to 5,000 users. Once set up, AD Connector offers the following benefits:
• Your end users and IT administrators can use their existing corporate credentials to log on to AWS applications such as Amazon Workspaces, Amazon WorkDocs, or Amazon WorkMail.
• You can manage AWS resources like Amazon EC2 instances or Amazon S3 buckets through 1AM role-based access to the AWS Management Console.
• You can consistently enforce existing security policies (such as password expiration, password history, and account lockouts) whether users or IT administrators are accessing resources in your on- premises infrastructure or in the AWS Cloud.
• You can use AD Connector to enable multi-factor authentication by integrating with your existing RADIUS-based MFA infrastructure to provide an additional layer of security when users access AWS applications.
For more information on the AD Connector, please visit the below URL:
• http://docs.aws.amazon.com/directoryservice/latest/admin-guide/directory_ad_con nector.htm I

NEW QUESTION 71
You need to deploy an AWS stack in a repeatable manner across multiple environments. You have selected CloudFormation as the right tool to accomplish this, but have found that there is a resource type you need to create and model, but is unsupported by CloudFormation. How should you overcome this challenge?

A. Use a CloudFormation Custom Resource Template by selecting an API call to proxy for create, update, and delete action
B. CloudFormation will use the AWS SDK, CLI, or API method of your choosing as the state transition function for the resource type you are modeling.
C. Submit a ticket to the AWS Forum
D. AWS extends CloudFormation Resource Types by releasing tooling to the AWS Labs organization on GitHu
E. Their response time is usually 1 day, and theycomplete requests within a week or two.
F. Instead of depending on CloudFormation, use Chef, Puppet, or Ansible to author Heat templates, which are declarative stack resource definitions that operate over the OpenStack hypervisor and cloud environment.
G. Create a CloudFormation Custom Resource Type by implementing create, update, and delete functionality, either by subscribing a Custom Resource Provider to an SNS topic, or by implementing the logic in AWS Lambda.

Answer: D

Explanation:
Custom resources enable you to write custom provisioning logic in templates that AWS Cloud Formation runs anytime you create, update (if you changed the custom resource), or delete stacks. For example, you might want to include resources that aren't available as AWS Cloud Formation resource types. You can include those resources by using custom resources. That way you can still manage all your related resources in a single stack.
Use the AWS:: Cloud Formation:: Custom Resource or Custom ::String resource type to define custom resources in your templates. Custom resources require one property: the service token, which specifies where AWS CloudFormation sends requests to, such as an Amazon SNS topic.
For more information on Custom Resources in Cloudformation, please visit the below U RL:
http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/template-custom- resources.html

NEW QUESTION 72
Your API requires the ability to stay online during AWS regional failures. Your API does not store any state, it only aggregates data from other sources - you do not have a database. What is a simple but effective way to achieve this uptime goal?

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

Answer: D

Explanation:
Failover routing lets you route traffic to a resource when the resource is healthy or to a different resource when the first resource is unhealthy. The primary and secondary resource record sets can route traffic to anything from an Amazon S3 bucket that is configured as a website to a complex tree of records.
For more information on Route53 Failover Routing, please visit the below URL:
http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html

NEW QUESTION 74
Your CTO thinks your AWS account was hacked. What is the only way to know for certain if there was unauthorized access and what they did, assuming your hackers are very sophisticated AWS engineers and doing everything they can to cover their tracks?

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

Answer: A

Explanation:
To determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it, you can use CloudTrail log file integrity validation. This feature is built using industry standard algorithms: SHA-256 for hashing and SHA-256 with RSA for digital signing. This makes it computationally infeasible to modify, delete or forge CloudTrail log files without detection. You can use the AWS CLI to validate the files in the location where CloudTrail delivered them
Validated log files are invaluable in security and forensic investigations. For example, a validated log file enables you to assert positively that the log file itself has not changed, or that particular user credentials performed specific API activity. The CloudTrail log file integrity validation process also lets you know if a log file has been deleted or changed, or assert positively that no log files were delivered to your account during a given period of time.
For more information on Cloudtrail log file validation, please visit the below URL:
http://docs.aws.a mazon.com/awscloudtrail/latest/userguide/cloudtrai l-log-file-validation- intro.html

NEW QUESTION 77
Your development team is using access keys to develop an application that has access to S3 and DynamoDB. A new security policy has outlined that the credentials should not be older than 2 months, and should be rotated. How can you achieve this

A. Use the application to rotate the keys in every 2 months via the SDK
B. Use a script which will query the date the keys are create
C. If older than 2 months, delete them and recreate new keys
D. Delete the user associated with the keys after every 2 month
E. Then recreate the user again.D- Delete the I AM Role associated with the keys after every 2 month
F. Then recreate the I AM Roleagain.

Answer: B

Explanation:
One can use the CLI command list-access-keys to get the access keys. This command also returns the "CreateDate" of the keys. If the CreateDate is older than 2 months, then the keys can be deleted.
The Returns list-access-keys CLI command returns information about the access key IDs associated with the specified I AM user. If there are none, the action returns
an empty list.
For more information on the CLI command, please refer to the below link: http://docs.aws.amazon.com/cli/latest/reference/iam/list-access-keys.html

NEW QUESTION 78
When creating an Elastic Beanstalk environment using the Wizard, what are the 3 configuration options presented to you

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

Answer: ABD

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

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

NEW QUESTION 81
You have carried out a deployment using Elastic Beanstalk with All at once method, but the application is unavailable. What could be the reason for this

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

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

Answer: C

Explanation:
The AWS Documentation mentions
Because Elastic Beanstalk uses a drop-in upgrade process, there might be a few seconds of downtime. Use rolling deployments to minimize the effect of deployments on your production environments.
For more information on troubleshooting Elastic Beanstalk, please refer to the below link:
• http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/troubleshooting-deployments.html
• https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.de ploy-existing- version, html

NEW QUESTION 82
Which of the below services can be used to deploy application code content stored in Amazon S3 buckets, GitHub repositories, or Bitbucket repositories

A. CodeCommit
B. CodeDeploy
C. S3Lifecycles
D. Route53

Answer: B

Explanation:
The AWS documentation mentions
AWS CodeDeploy is a deployment service that automates application deployments to Amazon EC2 instances or on-premises instances in your own facility.
For more information on Code Deploy please refer to the below link:
• http://docs.ws.amazon.com/codedeploy/latest/userguide/welcome.html

NEW QUESTION 85
When your application is loaded onto an Opsworks stack, which of the following event is triggered by Opsworks?

A. Deploy
B. Setup
C. Configure
D. Shutdown

Answer: A

Explanation:
When you deploy an application, AWS Ops Works Stacks triggers a Deploy event, which runs each layer's Deploy recipes. AWS OpsWorks Stacks also installs stack configuration and deployment attributes that contain all of the information needed to deploy the app, such as the app's repository and database connection data. For more information on the Deploy event please refer to the below link:
• http://docs.aws.amazon.com/opsworks/latest/userguide/workingapps.html

NEW QUESTION 89
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 94
You have the requirement to get a snapshot of the current configuration of the resources in your AWS Account. Which of the following services can be used for this purpose

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

Answer: C

Explanation:
The AWS Documentation mentions the following With AWS Config, you can do the following:
• Evaluate your AWS resource configurations for desired settings.
• Get a snapshot of the current configurations of the supported resources that are associated with your AWS account.
• Retrieve configurations of one or more resources that exist in your account.

• Retrieve historical configurations of one or more resources.
• Receive a notification whenever a resource is created, modified, or deleted.
• View relationships between resources. For example, you might want to find all resources that use a particular security group. For more information on AWS Config, please visit the below URL: http://docs.aws.amazon.com/config/latest/developerguide/WhatIsConfig.html

NEW QUESTION 97
You currently have an Autoscalinggroup that has the following settings Min capacity-2
Desired capacity - 2 Maximum capacity - 2
Your launch configuration has AMI'S which are based on the t2.micro instance type. The application running on these instances are now experiencing issues and you have identified that the solution is to change the instance type of the instances running in the Autoscaling Group.
Which of the below solutions will meet this demand.

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

Answer: C

Explanation:
You should make a copy of the launch configuration, add the new instance type. The change the Autoscaling Group to include the new instance type. Then change the Desired number of the Autoscaling Group to 4 so that instances of new instance type can be launched. Once launched, change the desired size back to 2, so that Autoscaling will delete the instances with the older configuration. Note that the assumption here is that the current instances are equally distributed across multiple AZ's because Autoscaling will first use the AZRebalance process to terminate instances.
Option A is invalid because you cannot make changes to an existing Launch configuration.
Option B is invalid because if you delete the existing launch configuration, then your application will not be available. You need to ensure a smooth deployment process.
Option D is invalid because you should change the desired size to 4 after attaching the new launch configuration.
For more information on Autoscaling Suspend and Resume, please visit the below URL: http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resu me-processes.html

NEW QUESTION 99
Your company has an on-premise Active Directory setup in place. The company has extended their footprint on AWS, but still want to have the ability to use their on-premise Active Directory for authentication. Which of the following AWS services can be used to ensure that AWS resources such as AWS Workspaces can continue to use the existing credentials stored in the on-premise Active Directory.

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

Answer: C

Explanation:
The AWS Documentation mentions the following
AD Connector is a directory gateway with which you can redirect directory requests to your on- premises Microsoft Active Directory without caching any information in the cloud. AD Connector comes in two sizes, small and large. A small AD Connector is designed for
smaller organizations of up to 500 users. A large AD Connector can support larger organizations of up to 5,000 users.
For more information on the AD connector, please refer to the below URL: http://docs.aws.amazon.com/directoryservice/latest/admin-guide/directory_ad_connector.html

NEW QUESTION 104
A company has developed a Ruby on Rails content management platform. Currently, OpsWorks with several stacks for dev, staging, and production is being used to deploy and manage the application. Now the company wants to start using Python instead of Ruby. How should the company manage the new deployment? Choose the correct answer from the options below

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

Answer: B

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

Please find the below link on a white paper for blue green deployments
• https://d03wsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 105
You are in charge of designing Cloudformation templates for your company. One of the key requirements is to ensure that if a Cloudformation stack is deleted, a snapshot of the relational database is created which is part of the stack. How can you achieve this in the best possible way?

A. Create a snapshot of the relational database beforehand so that when the cloudformation stack is deleted, the snapshot of the database will be present.
B. Use the Update policy of the cloudformation template to ensure a snapshot is created of the relational database.
C. Use the Deletion policy of the cloudformation template to ensure a snapshot is created of the relational database.
D. Create a new cloudformation template to create a snapshot of the relational database.

Answer: C

Explanation:
The AWS documentation mentions the following
With the Deletion Policy attribute you can preserve or (in some cases) backup a resource when its stack is deleted. You specify a DeletionPolicy attribute for each resource that you want to control. If a resource has no DeletionPolicy attribute, AWS Cloud Formation deletes the resource by default. Note that this capability also applies to update operations that lead to resources being removed.
For more information on the Deletion policy, please visit the below URL: http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/aws-attri bute-deletionpolicy.html

NEW QUESTION 106
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 108
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.htmI

NEW QUESTION 113
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/mon itor-charges.html

NEW QUESTION 115
Your IT company is currently hosting a production environment in Elastic beanstalk. You understand that the Elastic beanstalk service provides a facility known as Managed updates which are minor and patch version updates which are periodically required for your system. Your IT supervisor is worried about the impact that these updates would have on the system. What can you tell about the Elastic beanstalk service with regards to managed updates

A. Package updates can be configurable weekly maintenance window
B. Elastic Beanstalk applies managed updates with no downtime
C. Elastic Beanstalk applies managed updates with no reduction in capacity
D. All of the above

Answer: D

Explanation:
The AWS Documentation mentions the following on package updates for the Clastic beanstalk environment
You can configure your environment to apply minor and patch version updates automatically during a configurable weekly maintenance window with Managed Platform Updates. Elastic Beanstalk applies managed updates with no downtime or reduction in capacity, and cancels the update immediately if instances running your application on the new version fail health checks.
For more information on Elastic beanstalk managed updates please refer to the URL: https://docs.aws.a mazon.com/elasticbeanstalk/latest/dg/environment-platform -update-managed, html
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.platform.upgrade.html

NEW QUESTION 119
Of the 6 available sections on a Cloud Formation template (Template Description Declaration, Template Format Version Declaration, Parameters, Resources, Mappings, Outputs), which is the only one required for a CloudFormation template to be accepted? Choose an answer from the options below

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

Answer: D

Explanation:
If you refer to the documentation, you will see that Resources is the only mandatory field
Specifies the stack resources and their properties, such as an Amazon Elastic Compute Cloud instance or an Amazon Simple Storage Service bucket.
For more information on cloudformation templates, please refer to the below link:
• http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/template-anatomy.html

NEW QUESTION 120
Which of the following are components of the AWS Data Pipeline service. Choose 2 answers from the options given below

A. Pipeline definition
B. Task Runner
C. Task History
D. Workflow Runner

Answer: AB

Explanation:
The AWS Documentation mentions the following on AWS Pipeline
The following components of AWS Data Pipeline work together to manage your data: A pipeline definition specifies the business logic of your data management. A pipeline schedules and runs tasks. You upload your pipeline definition to the pipeline, and then activate the pipeline. You can edit the pipeline definition for a running pipeline and activate the pipeline again for it to take effect. You can deactivate the pipeline, modify a data source, and then activate the pipeline again. When you are finished with your pipeline, you can delete it.
Task Runner polls for tasks and then performs those tasks. For example. Task Runner could copy log files to Amazon S3 and launch Amazon EMR clusters. Task Runner is installed and runs automatically on resources created by your pipeline definitions. You can write a custom task runner application, or you can use the Task Runner application that is provided by AWS Data Pipeline.
For more information on AWS Pipeline, please visit the link: http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/what-is-datapipeline.html

NEW QUESTION 121
You were just hired as a DevOps Engineer for a startup. Your startup uses AWS for 100% of their infrastructure. They currently have no automation at all for deployment, and they have had many failures while trying to deploy to production.The company has told you deployment process risk mitigation is the most important thing now, and you have a lot of budget for tools and AWS resources.
Their stack includes a 2-tier API with data stored in DynamoDB or S3, depending on type. The Compute layer is EC2 in Auto Scaling Groups. They use Route53 for DNS pointing to an ELB. An ELB balances load across the EC2 instances. The scaling group properly varies between 4 and 12 EC2 servers. Which of the following approaches, given this company's stack and their priorities, best meets the company's needs?

A. Model the stack in AWS Elastic Beanstalk as a single Application with multiple Environment
B. Use Elastic Beanstalk's Rolling Deploy option to progressively roll out application code changes when promoting across environments.
C. Model the stack in three CloudFormation templates: Data layer, compute layer, and networking laye
D. Write stack deployment and integration testing automation following Blue-Green methodologie
E. •>/
F. Model the stack in AWS OpsWorks as a single Stack, with 1 compute layer and its associated EL
G. Use Chef and App Deployments to automate Rolling Deployment.
H. Model the stack in 1 CloudFormation template, to ensure consistency and dependency graph resolutio

I. Write deployment and integration testingautomation following Rolling Deployment methodologies.

Answer: B

Explanation:
Here you are using 2 of the best practices for deployment, one is Blue Green Deployments and the other is using Nested Cloudformation stacks.
The AWS Documentation mentions the below on nested stacks
As your infrastructure grows, common patterns can emerge in which you declare the same components in each of your templates. You can separate out these common components and create dedicated templates for them. That way, you can mix and match different templates but use nested stacks to create a single, unified stack. Nested stacks are stacks that create other stacks. To create nested stacks, use the AWS::CloudFormation::Stackresource in your template to reference other templates.
For more information on Cloudformation best practises, please visit the link:
• http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/best-practices. html For more information on Blue Green Deployment, please visit the link:
• https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 123
Which of the following are Lifecycle events available in Opswork? Choose 3 answers from the options below

A. Setup
B. Decommision
C. Deploy
D. Shutdown

Answer: ACD

Explanation:
Below is a snapshot of the Lifecycle events in Opswork.

For more information on Lifecycle events, please refer to the below URL:
• http://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-events.html

NEW QUESTION 127
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 stackcreation will wait until one service is created and launchedB.The user can use theHoldCondition resource to wait for the creation of the other dependent resources
B. The user can use theDependentCondition resource to hold the creation of the other dependentresources
C. The user can use the WaitConditionresource 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^ws.amazon.com/AWSCIoudFormation/latest/UserGuide/aws-properties-waitcondition.htmI

NEW QUESTION 128
A gaming company adopted AWS Cloud Formation to automate load-testing of theirgames. They have created an AWS Cloud Formation template for each gaming environment and one for the load- testing stack. The load-testing stack creates an Amazon Relational Database Service (RDS) Postgres database and two web servers running on Amazon Elastic Compute Cloud (EC2) that send HTTP requests, measure response times, and write the results into the database. A test run usually takes between 15 and 30 minutes. Once the tests are done, the AWS Cloud Formation stacks are torn down immediately. The test results written to the Amazon RDS database must remain accessible for visualization and analysis.
Select possible solutions that allow access to the test results after the AWS Cloud Formation load - testing stack is deleted.
Choose 2 answers.

A. Define an Amazon RDS Read-Replica in theload-testing AWS Cloud Formation stack and define a dependency relation betweenmaster and replica via the Depends On attribute.

B. Define an update policy to prevent deletionof the Amazon RDS database after the AWS Cloud Formation stack is deleted.
C. Define a deletion policy of type Retain forthe Amazon RDS resource to assure that the RDS database is not deleted with theAWS Cloud Formation stack.
D. Define a deletion policy of type Snapshotfor the Amazon RDS resource to assure that the RDS database can be restoredafter the AWS Cloud Formation stack is deleted.
E. Defineautomated backups with a backup retention period of 30 days for the Amazon RDSdatabase and perform point-in-time recovery of the database after the AWS CloudFormation stack is deleted.

Answer: CD

Explanation:
With the Deletion Policy attribute you can preserve or (in some cases) backup a resource when its stack is deleted. You specify a DeletionPolicy attribute for each resource that you want to control. If a resource has no DeletionPolicy attribute, AWS Cloud Formation deletes the resource by default.
To keep a resource when its stack is deleted, specify Retain for that resource. You can use retain for any resource. For example, you can retain a nested stack, S3 bucket, or CC2 instance so that you can continue to use or modify those resources after you delete their stacks.
For more information on Deletion policy, please visit the below url http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/aws-attri bute-deletionpolicy.html

NEW QUESTION 132
You are writing an AWS Cloud Formation template and you want to assign values to properties that will not be available until runtime. You know that you can use intrinsic functions to do this but are unsure as to which part of the template they can be used in. Which of the following is correct in describing how you can currently use intrinsic functions in an AWS CloudFormation template?

A. Youcan use intrinsic functions in any part of a template.
B. Youcan only use intrinsic functions in specific parts of a templat
C. You can useintrinsic functions in resource properties, metadata attributes, and updatepolicy attributes.
D. Youcan use intrinsic functions only in the resource properties part of a template.
E. Youcanuse intrinsic functions in any part of a template, exceptAWSTemplateFormatVersion and Description.

Answer: B

Explanation:
This is clearly given in the aws documentation. Intrinsic Function Reference
AWS Cloud Formation provides several built-in functions that help you manage your stacks. Use intrinsic functions in your templates to assign values to properties that are not available until runtime. Note
You can use intrinsic functions only in specific parts of a template. Currently, you can use intrinsic functions in resource properties, outputs, metadata attributes, and update policy attributes. You can also use intrinsic functions to conditionally create stack resources. For more information on intrinsic function please refer to the below link https://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/intrinsic-function -reference, html

NEW QUESTION 136
Which of the following run command types are available for opswork stacks? Choose 3 answers from the options given below.

A. UpdateCustom Cookbooks
B. Execute Recipes
C. Configure
D. UnDeploy

Answer: ABC

NEW QUESTION 139
Which of the following CLI commands is used to spin up new EC2 Instances?

A. awsec2 run-instances
B. awsec2 create-instances
C. awsec2 new-instancesD- awsec2 launch-instances

Answer: A

Explanation:
The AWS Documentation mentions the following
Launches the specified number of instances using an AMI for which you have permissions. You can specify a number of options, or leave the default options. The following rules apply:
[EC2-VPC] If you don't specify a subnet ID. we choose a default subnet from your default VPC for you. If you don't have a default VPC, you must specify a subnet ID in the request.
[EC2-Classic] If don't specify an Availability Zone, we choose one for you.
Some instance types must be launched into a VPC. if you do not have a default VPC. or if you do not specify a subnet ID. the request fails. For more information, see Instance Types Available Only in a VPC.
[EC2-VPC] All instances have a network interface with a primary private IPv4 address. If you don't specify this address, we choose one from the IPv4 range of your subnet.
Not all instance types support IPv6 addresses. For more information, see Instance Types.
If you don't specify a security group ID, we use the default security group. For more information, see Security Groups.
If any of the AMIs have a product code attached for which the user has not subscribed, the request fails. For more information on the Cc2 run instance command please refer to the below link http://docs.aws.a mazon.com/cli/latest/reference/ec2/run-instances.html

NEW QUESTION 142
Your company owns multiple AWS accounts. There is currently one development and one production account. You need to grant access to the development team to an S3 bucket in the production account. How can you achieve this?

A. Createan 1AM user in the Production account that allows users from the Developmentaccount (the trusted account) to access the S3 bucket in the Productionaccount.

B. When creating the role, define the Development account as a trustedentity and specify a permissions policy that allows trusted users to update theS3 bucket.
C. Use web identity federation with a third-partyidentity provider with AWS STS to grant temporary credentials and membershipinto the production 1AM user.
D. Createan 1AM cross account role in the Production account that allows users from theDevelopment account to access the S3 bucket in the Production account.

Answer: D

Explanation:
The AWS Documentation mentions the following on cross account roles
You can use AWS Identity and Access Management (1AM) roles and AWS Security Token Service (STS) to set up cross-account access between AWS accounts. When you assume an 1AM role in another AWS account to obtain cross-account access to services and resources in that account, AWS CloudTrail logs the cross-account activity. For more information on Cross account roles, please visit the below URL
• http://docs.aws.a mazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.htm l

NEW QUESTION 143
Which of the below 3 things can you achieve with the Cloudwatch logs service? Choose 3 options.

A. RecordAPI calls for your AWS account and delivers log files containing API calls toyour Amazon S3 bucket
B. Sendthe log data to AWS Lambda for custom processing or to load into other systems
C. Streamthe log data to Amazon Kinesis
D. Streamthe log data into Amazon Elasticsearch in near real-time with Cloud Watch Logssubscriptions.

Answer: BCD

Explanation:
You can use Amazon CloudWatch Logs to monitor, store, and access your log files from Amazon Elastic Compute Cloud (Amazon L~C2) instances, AWS CloudTrail, and other sources. You can then retrieve the associated log data from CloudWatch Logs.
For more information on Cloudwatch logs, please visit the below URL http://docs.ws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html

NEW QUESTION 144
Your firm has uploaded a large amount of aerial image data to S3. In the past, in your on-premises environment, you used a dedicated group of servers to process this data and used Rabbit MQ - 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 topassjob 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 147
Your company has recently extended its datacenter into a VPC on AWS. There is a requirement for on-premise users manage AWS resources from the AWS console. You don't want to create 1AM users for them again. Which of the below options will fit your needs for authentication?

A. UseOAuth 2.0 to retrieve temporary AWS security credentials to enable your membersto sign in to the AWS Management Console.
B. Useweb Identity Federation to retrieve AWS temporary security credentials toenable your members to sign in to the AWS Management Console.
C. Useyour on-premises SAML 2 O-compliant identity provider (IDP) to grant themembers federated access to the AWS Management Console via the AWS singlesign-on (SSO) endpoint.
D. Useyour on-premises SAML2.0-compliant identity provider (IDP) to retrieve temporarysecurity credentials to enable members to sign in to the AWS ManagementConsole.

Answer: C

Explanation:
You can use a role to configure your SAML 2.0-compliant IdP and AWS to permit your federated users to access the AWS Management Console. The role grants the user permissions to carry out tasks in the console.
For more information on aws SAML, please visit the below URL
• http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_ena ble-console- saml.html

NEW QUESTION 152
Explain what the following resource in a CloudFormation template does? Choose the best possible answer.

A. Createsan SNS topic which allows SQS subscription endpoints to be added as a parameteron thetemplate
B. Createsan SNS topic that allow SQS subscription endpoints
C. Createsan SNS topic and then invokes the call to create an SQS queue with a logicalresource name of SQSQueue
D. Creates an SNS topic and adds asubscription ARN endpoint for the SQS resource created under the logical nameSQSQueue

Answer: D

Explanation:
The intrinsic function Fn::GetAtt returns the value of an attribute from a resource in the template. This has nothing to do with adding parameters (Option A is wrong) or allowing endpoints (Option B is wrong) or invoking relevant calls (Option C is wrong)
For more information on Fn:: GetAtt function please refer to the below link
http://docs.aws.a mazon.com/AWSCIoudFormation/latest/UserGuide/intrinsic-function -reference- getatt.htm I

NEW QUESTION 157
What would you set in your CloudFormation template to fire up different instance sizes based off of environment type? i.e. (If this is for prod, use m1.large instead of t1.micro)

A. Outputs
B. Resources
C. Mappings
D. conditions

Answer: D

Explanation:
The optional Conditions section includes statements that define when a resource is created or when a property is defined. For example, you can compare whether a value is equal to another value. Based on the result of that condition, you can conditionally create resources. If you have multiple conditions, separate them with commas.
For more information on Cloudformation conditions please visit the below link
http://docs^ws.a mazon.com/AWSCIoudFormation/latest/UserGuide/cond itions-section- structure.htm I

NEW QUESTION 158
How can you resolve a dependency Error when using Cloudformation?

A. Usethe mappings attribute
B. Usethe parameter attribute
C. UsetheDependsOn 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 Cloud Formation 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/AWSCIoudFormation/latest/UserGuide/troubleshooting.html

NEW QUESTION 163
You are in charge of designing a Cloudformation template which deploys a LAMP stack. After deploying a stack, you see that the status of the stack is showing as CREATE_COMPLETE, but the apache server is still not up and running and is experiencing issues while starting up. You want to ensure that the stack creation only shows the status of CREATE_COMPLETE after all resources defined in the stack are up and running. How can you achieve this?
Choose 2 answers from the options given below.

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

Answer: CD

Explanation:
The AWS Documentation mentions
When you provision an Amazon EC2 instance in an AWS Cloud Formation stack, you might specify additional actions to configure the instance, such as install software packages or bootstrap applications. Normally, CloudFormation proceeds with stack creation after the instance has been successfully created. However, you can use a Creation Pol icy so that CloudFormation proceeds with stack creation only after your configuration actions are done. That way you'll know your applications are ready to go after stack creation succeeds.
For more information on the Creation Policy, please visit the below url https://aws.amazon.com/blogs/devops/use-a-creationpolicy-to-wait-for-on-instance-configurations/

NEW QUESTION 164
You are in charge of designing a number of Cloudformation templates for your organization. You are required to make changes to stack resources every now and then based on the requirement. How can you check the impact of the change to resources in a cloudformation stack before deploying changes to the stack?

A. Thereis no way to control thi
B. You need to check for the impact beforehand.
C. UseCloudformation change sets to check for the impact to the changes.
D. UseCloudformation Stack Policies to check for the impact to the changes.

E. UseCloudformation Rolling Updates to check for the impact to the changes.

Answer: B

Explanation:
The AWS Documentation mentions
When you need to update a stack, understanding how your changes will affect running resources before you implement them can help you update stacks with confidence. Change sets allow you to preview how proposed changes to a stack might impact your running resources, for example, whether your changes will delete or replace any critical resources, AWS CloudFormation makes the changes to your stack only when you decide to execute the change set, allowing you to decide whether to proceed with your proposed changes or explore other changes by creating another change set. You can create and manage change sets using the AWS
CloudFormation console, AWS CLI, or AWS CloudFormation API.
For more information on Cloudformation change sets, please visit the below url http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks-changesets.html

NEW QUESTION 169
Which of the following is false when it comes to using the Elastic Load balancer with Opsworks stacks?

A. Youcan attach only one load balancer to a layer.
B. A Classic Load Balancer can span across AWSOpsWorks Stacks layers.
C. Eachload balancer can handle only one layer.
D. Youneed to create the load balancer before hand and then attach it to the Opsworkstack.

Answer: B

Explanation:
The AWS Documentation mentions the following
To use Clastic Load Balancing with a stack, you must first create one or more load balancers in the same region by using the Clastic Load Balancing console, CLI, or API. You should be aware of the following:
You can attach only one load balancer to a layer. Cach 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 Clastic Load Balancer with Opswork,
please visit the below url http://docs.aws.a mazon.com/opsworks/latest/userguide/layers-elb.html

NEW QUESTION 170
Which of the following are true with regard to Opsworks stack Instances? Choose 3 answers from the options given below.

A. Astacks instances can be a combination of both Linux and Windows based operatingsystems.
B. You can use EC2 Instances that were createdoutisde the boundary of Opswork.
C. Youcanuseinstancesrunningonyourownhardware.
D. Youcan start and stop instances manually.

Answer: BCD

Explanation:
The AWS Documentation mentions the following
1) You can start and stop instances manually or have AWS Ops Works Stacks automatically scale the number of instances. You can use time-based automatic scaling with any stack; Linux stacks also can use load-based scaling.
2) In addition to using AWS OpsWorks Stacks to create Amazon L~C2 instances, you can also register instances with a Linux stack that were created outside of AWS Ops Works Stacks. This includes Amazon CC2 instances and instances running on your own hardware. However, they must be running one of the supported Linux distributions. You cannot register Amazon CC2 or on-premises Windows instances.
3) A stack's instances can run either Linux or Windows. A stack can have different Linux versions or distributions on different instances, but you cannot mix Linux and Windows instances.
For more information on Opswork instances, please visit the below url http://docs.aws.a mazon.co m/o psworks/latest/usergu ide/workinginstances-os. html

NEW QUESTION 171
Which of the following tools for EC2 can be used to administer instances without the need to SSH or RDP into the instance.

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

Answer: C

Explanation:
You can use Run Command from the Amazon L~C2 console to configure instances without having to login to each instance
For more information on the Run Command, please visit the below URL:
• http://docs.aws.a mazon.com/systems-manager/latest/userguide/rc-console.html

NEW QUESTION 173
You are creating a Cloudformation template in which Userdata isgoing to be passed to underlying EC2 Instance. Which of the below functions is normally used to pass data to the UserData section in the Cloudformation template?

A. "UserData": { "Fn::Base64":{
B. "UserData": < "Fn::Ref*:{
C. "UserData": { "Fn::GetAtt":{
D. "UserData": { "Fn::FindlnMap":{

Answer: A

Explanation:
The AWS Documentation mentions
The intrinsic function Fn::Base64 returns the Base64 representation of the input string. This function is typically used to pass encoded data to Amazon EC2 instances by way of the User Data property.
For more information on the Fn::Base64 function, please visit the below URL: http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/intrinsic-function -reference- base64.htm I

NEW QUESTION 177
Which of the following Cloudformation helper scripts can help install packages on EC2 resources

A. cfn-init
B. cfn-signal
C. cfn-get-metadata
D. cfn-hup

Answer: A

Explanation:
The AWS Documentation mentions
Currently, AWS CloudFormation provides the following helpers:
cf n-init: Used to retrieve and interpret the resource metadata, installing packages, creating files and starting services.
cf n-signal: A simple wrapper to signal an AWS CloudFormation CreationPolicy or WaitCondition,
enabling you to synchronize other resources in the stack with the application being ready.
cf n-get-metadata: A wrapper script making it easy to retrieve either all metadata defined for a resource or path to a specific key or subtree of the resource metadata.
cf n-hup: A daemon to check for updates to metadata and execute custom hooks when the changes are detected. For more information on helper scripts, please visit the below URL: http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/cfn-helper-scripts-reference.htmI

NEW QUESTION 181
Which of the following are the basic stages of a CI/CD Pipeline. Choose 3 answers from the options below

A. SourceControl
B. Build
C. Run
D. Production

Answer: ABD

Explanation:
The below diagram shows the stages of a typical CI/CD pipeline

For more information on AWS Continuous Integration, please visit the below URL: https://da.wsstatic.com/whitepapers/DevOps/practicing-continuous-integration-continuous- delivery-on- AWS.pdf

NEW QUESTION 184
As part of your continuous deployment process, your application undergoes an I/O load performance test before it is deployed to production using new AMIs. The application uses one Amazon EBS PIOPS volume per instance and requires consistent I/O performance.
Which of the following must be carried out to ensure that I/O load performance tests yield the correct results in a repeatable manner?

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

Answer: B

Explanation:
Since the AMI will get all the data from S3 as snapshots, always ensure the volume prewarmed before it is set for the load test.
For more information on benchmarking procedures please see the below link:
• hrtp://docs^ws.amazon.com/AWSCC2/latest/UserGuide/berK;hmark_prooedures.html

NEW QUESTION 185
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. Publishyour data to CloudWatch Logs, and configure your application to autoscale tohandle the load on demand.
B. Publishyour log data to an Amazon S3 bucke
C. Use AWS CloudFormation to create an AutoScalinggroup to scale your post-processing application which is configured topull down your log files stored in Amazon S3.
D. Postyour log data to an Amazon Kinesis data stream, and subscribe yourlog-processing application so that is configured to process your logging data.
E. Configurean Auto Scalinggroup to increase the size of your Amazon EMR cluster

Answer: C

Explanation:
The AWS Documentation mentions the below
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.
Amazon Kinesis enables you to process and analyze data as it
arrives and respond in real-time instead of having to wait until all your data is collected before the processing can begin.
For more information on AWS Kinesis please see the below link:
• https://aws.amazon.com/kinesis/

NEW QUESTION 186
You are managing an application that contains Go as the front end, MongoDB for document management and is hosted on a relevant Web server. You pre-bake AMI'S with the latest version of the Web server, then user the User Data section to setup the application. You now have a change to the underlying Operating system version and need to deploy that accordingly. How can this be done
in the easiest way possible.

A. Createa new EBS Volume with the relevant OS patches and attach it to the EC2Instance.
B. Createa Cloudformation stack with the new AMI and then deploy the applicationaccordingly.
C. Createa new pre-baked AM I with the new OS and use the User Data seciton to deploy theapplication.
D. Createan Opswork stack with the new AMI and then deploy the application accordingly.

Answer: C

Explanation:
The best way in this scenario is to continue the same deployment process which was being used and create a new AMI and then use the User Data section to deploy the application.
For more information on AWS AMI's please see the below link:
• http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/AMIs.htmls

NEW QUESTION 188
You have an Autoscaling Group which is launching a set of t2.small instances. You now need to replace those instances with a larger instance type. How would you go about making this change in an ideal manner?

A. Changethe Instance type in the current launch configuration to the new instance type.
B. Createanother Autoscaling Group and attach the new instance type.
C. Createa new launch configuration with the new instance type and update yourAutoscaling Group.
D. Changethe Instance type of the Underlying EC2 instance directly.

Answer: C

Explanation:
Answer - C
The AWS Documentation mentions
A launch configuration is a template that an Auto Scaling group uses to launch EC2 instances. When you create a launch configuration, you specify information for the instances such as the ID of the Amazon Machine Image (AMI), the instance type, a key pair, one or more security groups, and a block device mapping. If you've launched an EC2 instance before, you specified the same information in order to launch the instance. When you create an Auto Scalinggroup, you must specify a launch configuration. You can specify your launch configuration with multiple Auto Scaling groups.
However, you can only specify one launch configuration for an Auto Scalinggroup at a time, and you can't modify a launch configuration after you've created it. Therefore, if you want to change the launch configuration for your Auto Scalinggroup, you must create a launch configuration and then update your Auto Scaling group with the new launch configuration.
For more information on launch configurations please see the below link:
• http://docs.aws.amazon.com/autoscaling/latest/userguide/l_au nchConfiguration.html

NEW QUESTION 190
You are planning on configuring logs for your Elastic Load balancer. At what intervals does the logs get produced by the Elastic Load balancer service. Choose 2 answers from the options given below

A. 5minutes
B. 60minutes
C. 1 minute
D. 30seconds

Answer: AB

Explanation:
The AWS Documentation mentions
Clastic Load Balancing publishes a log file for each load balancer node at the interval you specify. You can specify a publishing interval of either 5 minutes or 60 minutes when you enable the access log for your load balancer. By default. Elastic Load Balancing publishes logs at a 60-minute interval.
For more information on Elastic load balancer logs please see the below link: http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/access-log-collection.html

NEW QUESTION 191
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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