DOP-C01 Dumps

AWS Certified DevOps Engineer- Professional

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

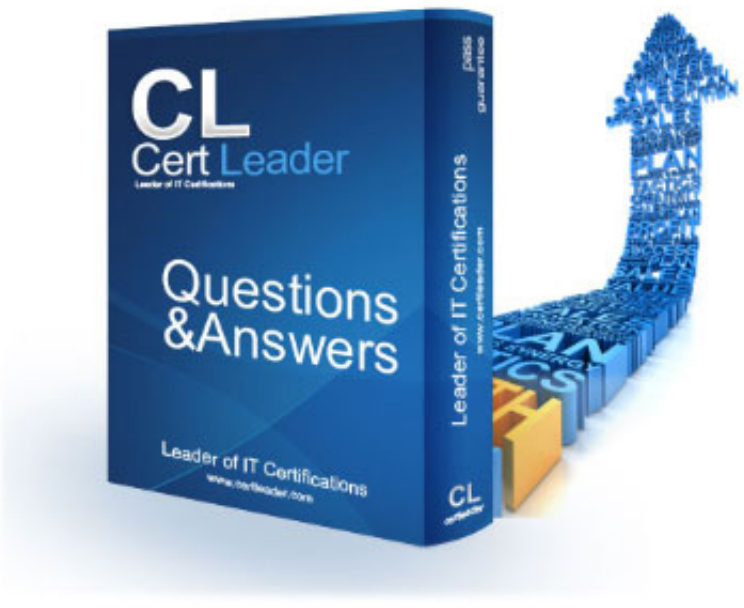

NEW QUESTION 1
You have an application running a specific process that is critical to the application's functionality, and have added the health check process to your Auto Scaling Group. The instances are showing healthy but the application itself is not working as it should. What could be the issue 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 2
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 3
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.htmI

NEW QUESTION 4
You currently run your infrastructure on Amazon EC2 instances behind an Auto Scalinggroup. All logs for your application are currentl\ written to ephemeral

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

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

Answer: C

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

NEW QUESTION 5
You have a code repository that uses Amazon S3 as a data store. During a recent audit of your security controls, some concerns were raised about maintaining the integrity of the data in the Amazon S3 bucket. Another concern was raised around securely deploying code from Amazon S3 to applications running on Amazon EC2 in a virtual private cloud. What are some measures that you can implement to mitigate these concerns? Choose two answers from the options given below.

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

Answer: BD

Explanation:
You can add another layer of protection by enabling MFA Delete on a versioned bucket. Once you do
so, you must provide your AWS account's access keys and a
valid code from the account's MFA device in order to permanently delete an object version or suspend or reactivate versioning on the bucket.
For more information on MFA please refer to the below link: https://aws.amazon.com/blogs/security/securing-access-to-aws-using-mfa-part-3/
IAM roles are designed so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that the applications use. Instead of creating and distributing your AWS credentials, you can delegate permission to make API requests using 1AM roles For more information on Roles for CC2 please refer to the below link: http://docs.aws.a mazon.com/AWSCC2/latest/UserGuide/iam-roles-for-amazon-ec2. html
Option A is invalid because this will not address either the integrity or security concern completely. Option C is invalid because user credentials should never be used in CC2 instances to access AWS resources.
Option C and F are invalid because AWS Pipeline is an unnecessary overhead when you already have inbuilt controls to manager security for S3.

NEW QUESTION 6
The operations team and the development team want a single place to view both operating system and application logs. How should you implement this using A WS services? Choose two from the options below

A. Using AWS CloudFormation, create a Cloud Watch Logs LogGroup and send the operating system and application logs of interest using the Cloud Watch Logs Agent.
B. Using AWS CloudFormation and configuration management, set up remote logging to send events via UDP packets to CloudTrail.
C. Using configuration management, set up remote logging to send events to Amazon Kinesis and insert these into Amazon CloudSearch or Amazon Redshift, depending on available analytic tools.
D. Using AWS CloudFormation, merge the application logs with the operating system logs, and use 1AM Roles to allow both teams to have access to view console output from Amazon EC2.

Answer: AC

Explanation:
Option B is invalid because Cloudtrail is not designed specifically to take in UDP packets
Option D is invalid because there are already Cloudwatch logs available, so there is no need to have specific logs designed for this.
You can use Amazon CloudWatch Logs to monitor, store, and access your log files from Amazon Elastic Compute Cloud (Amazon 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 refer to the below link:
http://docs^ws.amazon.com/AmazonCloudWatch/latest/logs/WhatlsCloudWatchLogs.html You can the use Kinesis to process those logs
For more information on Amazon Kinesis please refer to the below link: http://docs.aws.a mazon.com/streams/latest/dev/introduction.html

NEW QUESTION 7
You have been requested to use CloudFormation to maintain version control and achieve automation for the applications in your organization. How can you best use CloudFormation to keep everything agile and maintain multiple environments while keeping cost down?

A. Create separate templates based on functionality, create nested stacks with CloudFormation.

B. Use CloudFormation custom resources to handle dependencies between stacks
C. Create multiple templates in one CloudFormation stack.
D. Combine all resources into one template for version control and automation.

Answer: A

Explanation:
As your infrastructure grows, common patterns can emerge in which you declare the same components in each of your templates. You can separate out these common components and create dedicated templates for them. That way, you can mix and match different templates but use nested stacks to create a single, unified stack. Nested stacks are stacks that create other stacks. To create nested stacks, use the AWS:: Cloud Form ation::Stackresource in your template to reference
other templates. For more information on Cloudformation best practises please refer to the below link:
http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/best-practices.html

NEW QUESTION 8
If your application performs operations or workflows that take a long time to complete, what service can the Elastic Beanstalk environment do for you?

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

Answer: A

Explanation:
Elastic Beanstalk simplifies this process by managing the Amazon SQS queue and running a daemon process on each instance that reads from the queue for you. When the daemon pulls an item from the queue, it sends an HTTP POST request locally to http://localhost/ with the contents of the queue message in the body. All that your application needs to do is perform the long-running task in response to the POST.
For more information Elastic Beanstalk managing worker environments, please visit the below URL:
? http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.htm I

NEW QUESTION 9
You have an Auto Scaling group with an Elastic Load Balancer. You decide to suspend the Auto Scaling AddToLoadBalancer for a short period of time. What will happen to the instances launched during the suspension period?

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

Answer: C

Explanation:
If you suspend AddTo Load Balancer, Auto Scaling launches the instances but does not add them to the load balancer or target group. If you resume the AddTo Load Balancer process. Auto Scaling resumes adding instances to the load balancer or target group when they are launched. However, Auto Scaling does
not add the instances that were launched while this process was suspended. You must register those instances manually.
For more information on the Suspension and Resumption process, please visit the below U RL: http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resume-processes.html

NEW QUESTION 10
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
You have a multi-docker environment that you want to deploy to AWS. Which of the following configuration files can be used to deploy a set of Docker containers as an Elastic Beanstalk application?

A. Dockerrun.awsjson
B. .ebextensions
C. Dockerrunjson
D. Dockerfile

Answer: A

Explanation:
A Dockerrun.aws.json file is an Clastic Beanstalk-specific JSON file that describes how to deploy a set of Docker containers as an Clastic Beanstalk application.
You can use aDockerrun.aws.json file for a multicontainer Docker environment.
Dockerrun.aws.json describes the containers to deploy to each container instance in the environment as well as the data volumes to create on the host instance for
the containers to mount. http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_v2config.html

NEW QUESTION 11
Your current log analysis application takes more than four hours to generate a report of the top 10 users of your web application. You have been asked to implement a system that can report this information in real time, ensure that the report is always up to date, and handle increases in the number of requests to your web application. Choose the option that is cost-effective and can fulfill the requirements.

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

Answer: C

Explanation:
When you see Amazon Kinesis as an option, this becomes the ideal option to process data in real time.
Amazon Kinesis makes it easy to collect, process, and analyze real-time, streaming data so you can get timely insights and react quickly to new information. Amazon
Kinesis offers key capabilities to cost effectively process streaming data at any scale, along with the flexibility to choose the tools that best suit the requirements of your application. With Amazon Kinesis, you can ingest real-time data such as application logs, website clickstreams, loT telemetry data, and more into your databases, data lakes and data warehouses, or build your own real-time applications using this data. For more information on Amazon Kinesis, please visit the below URL:
• https://aws.amazon.com/kinesis

NEW QUESTION 14
After reviewing the last quarter's monthly bills, management has noticed an increase in the overall bill from Amazon. After researching this increase in cost, you discovered that one of your new services is doing a lot of GET Bucket API calls to Amazon S3 to build a metadata cache of all objects in the applications bucket. Your boss has asked you to come up with a new cost-effective way to help reduce the amount of these new GET Bucket API calls. What process should you use to help mitigate the cost?

A. Update your Amazon S3 buckets' lifecycle policies to automatically push a list of objects to a new bucket, and use this list to view objects associated with the application's bucket.
B. Create a new DynamoDB tabl
C. Use the new DynamoDB table to store all metadata about all objects uploaded to Amazon S3. Any time a new object is uploaded, update the application's 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 17
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 19
You have a complex system that involves networking, 1AM policies, and multiple, three-tier applications. You are still receiving requirements for the new system, so you don't yet know how many AWS components will be present in the final design. You want to start using AWS CloudFormation to define these AWS resources so that you can automate and version-control your infrastructure. How would you use AWS CloudFormation to provide agile new environments for your customers in a cost-effective, reliable manner?

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

Answer: B

Explanation:
As your infrastructure grows, common patterns can emerge in which you declare the same components in each of your templates. You can separate out these common components and create dedicated templates for them. That way, you can mix and match different templates but use nested stacks to create a single, unified stack. Nested stacks are stacks that create other stacks. To create nested stacks, use the AWS::CloudFormation::Stackresource in your template to reference other templates.
For more information on Cloudformation best practises please refer to the below link: http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/best-practices.html

NEW QUESTION 23
When an Auto Scaling group is running in Amazon Elastic Compute Cloud (EC2), your application rapidly scales up and down in response to load within a 10-minute window; however, after the load peaks, you begin to see problems in your configuration management system where previously terminated Amazon EC2 resources are still showing as active. What would be a reliable and efficient way to handle the cleanup of Amazon EC2 resources within your configuration management system? Choose two answers from the options given below

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

Answer: AD

Explanation:
There is a rich brand of CLI commands available for Cc2 Instances. The CLI is located in the following link:
• http://docs.aws.a mazon.com/cli/latest/reference/ec2/
You can then use the describe instances command to describe the EC2 instances.
If you specify one or more instance I Ds, Amazon CC2 returns information for those instances. If you do not specify instance IDs, Amazon EC2 returns information for all relevant instances. If you specify an instance ID that is not valid, an error is returned. If you specify an instance that you do not own, it is not included in the returned results.
• http://docs.aws.a mazon.com/cli/latest/reference/ec2/describe-insta nces.html
You can use the CC2 instances to get those instances which need to be removed from the configuration management system.

NEW QUESTION 24
You have a web application that's developed in Node.js The code is hosted in Git repository. You want to now deploy this application to AWS. Which of the below 2 options can fulfil this requirement.

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

J. S
K. Create an AWS CloudFormation template which creates an instance with the AWS::EC2::Instance resource type and an AMI with Docker pre-installe
L. With UserData, install Git to download the Node.js application and then set it up.

Answer: CD

Explanation:
Option A is invalid because there is no "awsgitpush" command
Option B is invalid because there is no AWS::CC2::Container resource type.
Clastic Beanstalk supports the deployment of web applications from Docker containers. With Docker containers, you can define your own runtime environment. You can choose your own platform, programming language, and any application dependencies (such as package managers or tools), that aren't supported by other platforms. Docker containers are self-contained and include all the configuration information and software your web application requires to run.
For more information on Docker and Clastic beanstalk please refer to the below link:
? http://docs.aws.a mazon.com/elasticbeanstalk/latest/dg/create_deploy_docker.html
When you launch an instance in Amazon CC2, you have the option of passing user data to the instance that can be used to perform common automated configuration tasks and even run scripts after the instance starts. You can pass two types of user data to Amazon CC2: shell scripts and cloud- init directives. You can also pass this data into the launch wizard as plain text, as a file (this is useful for launching instances using the command line tools), or as base64-encoded text (for API calls). For more information on Cc2 User data please refer to the below link:
• http://docs.aws.a mazon.com/AWSCC2/latest/UserGuide/user-data. htm l
Note: "git aws.push" with CB CLI 2.x - see a forum thread at https://forums.aws.amazon.com/thread.jspa7messageID=583202#jive-message-582979. Basically, this is a predecessor to the newer "eb deploy" command in CB CLI 31. This question kept in order to be consistent with exam.

NEW QUESTION 26
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 l

NEW QUESTION 29
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 32
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 35
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

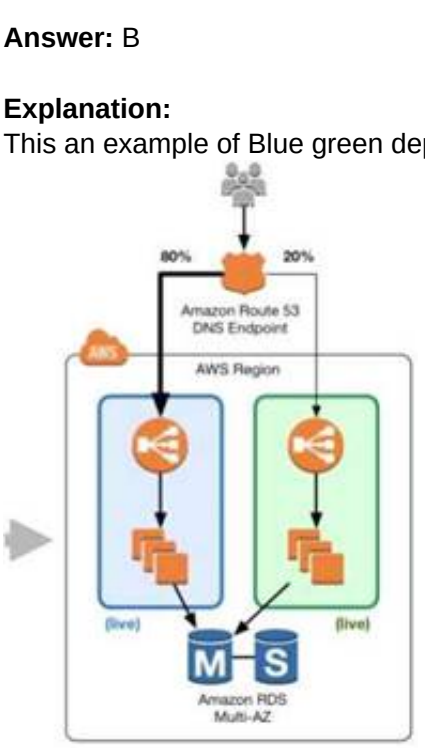

With Amazon Route 53, you can define a percentage of traffic to go to the green environment and gradually update the weights until the green environment carries the full production traffic. A weighted distribution provides the ability to perform canary analysis where a small percentage of production 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 40
You've been tasked with improving the current deployment process by making it easier to deploy and reducing the time it takes. You have been tasked with creating a continuous integration (CI) pipeline that can build AMI'S. Which of the below is the best manner to get this done. Assume that at max your development team will be deploying builds 5 times a week.

A. Use a dedicated EC2 instance with an EBS Volum
B. Download and configure the code and then crate an AMI out of that.
C. Use OpsWorks to launch an EBS-backed instance, then use a recipe to bootstrap the instance, and then have the CI system use the CreateImage API call to

make an AMI from it.
D. Upload the code and dependencies to Amazon S3, launch an instance, download the package 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 45
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/WhatIsCloud 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 50
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 55
You have decided to migrate your application to the cloud. You cannot afford any downtime. You want to gradually migrate so that you can test the application with a small percentage of users and increase over time. Which of these options should you implement?

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

Answer: D

Explanation:
Option A is incorrect because DirectConnect cannot control the flow of traffic.
Option B is incorrect because you want to split the percentage of traffic. Failover will direct all of the traffic to the backup servers.
Option C is incorrect because you cannot control the percentage distribution of traffic.
Weighted routing lets you associate multiple resources with a single domain name (example.com) or subdomain name (acme.example.com) and choose how much traffic is routed to each resource. This can be useful for a variety of purposes, including load
balancing and testing new versions of software.
For more information on the Routing policy please refer to the below link: http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html

NEW QUESTION 56
You are using Chef in your data center. Which service is designed to let the customer leverage existing Chef recipes in AWS?

A. AWS Elastic Beanstalk
B. AWSOpsWorks
C. AWS CloudFormation
D. Amazon Simple Workflow Service

Answer: B

Explanation:
AWS OpsWorks is a configuration management service that uses Chef, an automation platform that treats server configurations as code. OpsWorks uses Chef to automate how servers are configured, deployed, and managed across your Amazon Clastic Compute Cloud (Amazon CC2) instances or on-premises compute environments. OpsWorks has two offerings, AWS Opsworks for Chef Automate, and AWS OpsWorks Stacks.
For more information on Opswork and SNS please refer to the below link:
• https://aws.amazon.com/opsworks/

NEW QUESTION 57
You are creating a new API for video game scores. Reads are 100 times more common than writes, and the top 1% of scores are read 100 times more frequently than the rest of the scores. What's the best design for this system, using DynamoDB?

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

Answer: D

Explanation:
Because the lOOx read ratio is mostly driven by a small subset, with caching, only a roughly equal number of reads to writes will miss the cache, since the supermajority will hit the top 1% scores. Knowing we need to set the values roughly equal when using caching, we select AWS OastiCache, because CloudFront cannot directly cache DynamoDB queries, and OastiCache is an excellent in-memory cache for database queries, rather than a distributed proxy cache for content delivery.
For more information on DynamoDB table gudelines please refer to the below link:
• http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GuidelinesForTables.html

NEW QUESTION 60
You are planning on using the Amazon RDS facility for Fault tolerance for your application. How does Amazon RDSmuIti Availability Zone model work

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

Answer: A

Explanation:
Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB)
Instances, making them a natural fit for production database
workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). 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.
The below diagram from the AWS documentation shows how this is configured

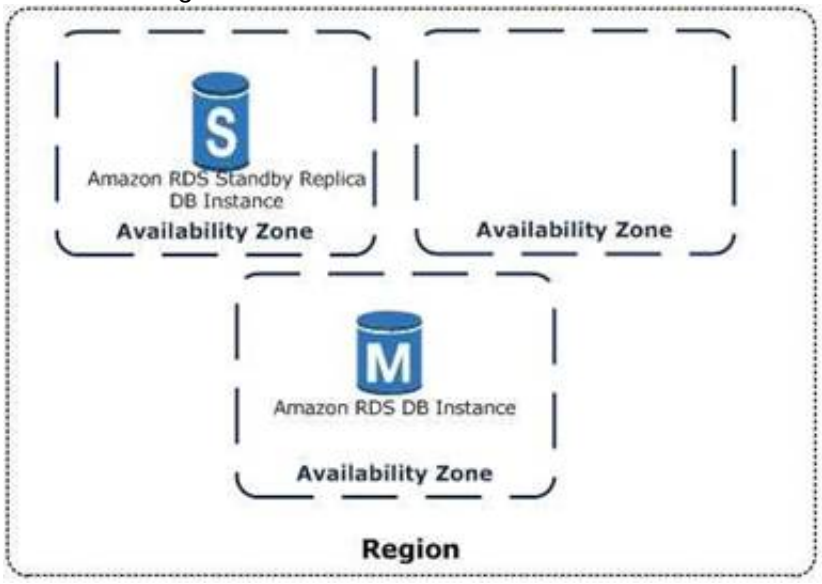

Option B is invalid because the replication is synchronous.
Option C and D are invalid because this is built around AZ and not regions. For more information on Multi-AZ RDS, please visit the below URL:
https://aws.amazon.com/rds/details/multi-az/

NEW QUESTION 65
There is a requirement to monitor API calls against your AWS account by different users and entities. There needs to be a history of those calls. The history of those calls are needed in in bulk for later review. Which 2 services can be used in this scenario

A. AWS Config; AWS Inspector
B. AWS CloudTrail; AWS Config
C. AWS CloudTrail; CloudWatch Events

D. AWS Config; AWS Lambda

Answer: C

Explanation:
You can use AWS CloudTrail to get a history of AWS API calls and related events for your account. This history includes calls made with the AWS Management Console, AWS Command Line Interface, AWS SDKs, and other AWS services. For more information on Cloudtrail, please visit the below URL:
• http://docs.aws.a mazon.com/awscloudtrail/latest/userguide/cloudtrai l-user-guide.html
Amazon Cloud Watch Cvents delivers a near real-time stream of system events that describe changes in Amazon Web Services (AWS) resources. Using simple rules that you can quickly set up, you can match events and route them to one or more target functions or streams. Cloud Watch Cvents becomes aware of operational changes as they occur. Cloud Watch Cvents responds to these operational changes and takes corrective action as necessary, by sending messages to respond to the environment, activating functions, making changes, and capturing state information. For more information on Cloud watch events, please visit the below U RL:
• http://docs.aws.a mazon.com/AmazonCloudWatch/latest/events/Whatl sCloudWatchCvents.html

NEW QUESTION 69
You want to pass queue messages that are 1GB each. How should you achieve this?

A. Use Kinesis as a buffer stream for message bodie
B. Store the checkpoint id for the placement in the Kinesis Stream in SQS.
C. Use the Amazon SQS Extended Client Library for Java and Amazon S3 as a storage mechanism for message bodies.
D. Use SQS's support for message partitioning and multi-part uploads on Amazon S3.
E. Use AWS EFS as a shared pool storage mediu
F. Store filesystem pointers to the files on disk in the SQS message bodies.

Answer: B

Explanation:
You can manage Amazon SQS messages with Amazon S3. This is especially useful for storing and consuming messages with a message size of up to 2 GB. To manage
Amazon SQS messages with Amazon S3, use the Amazon SQS Extended Client Library for Java. Specifically, you use this library to:
• Specify whether messages are always stored in Amazon S3 or only when a message's size exceeds 256 KB.
• Send a message that references a single message object stored in an Amazon S3 bucket.
• Get the corresponding message object from an Amazon S3 bucket.
• Delete the corresponding message object from an Amazon S3 bucket.
For more information on processing large messages for SQS, please visit the below URL:
http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-s3- messages. html

NEW QUESTION 72
For AWS Auto Scaling, what is the first transition state an existing instance enters after leaving Standby state?

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

Answer: C

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

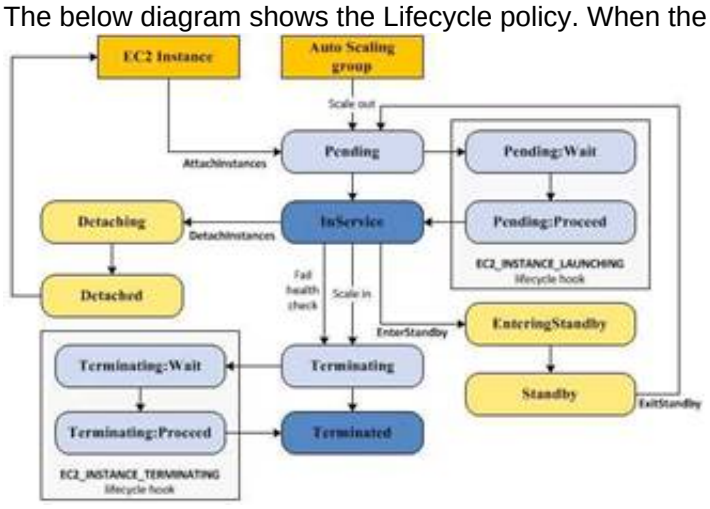

For more information on Autoscaling Lifecycle, please refer to the below link:
http://docs.aws.amazon.com/autoscaling/latest/userguide/AutoScaingGroupLifecycle.html

NEW QUESTION 75
You have deployed a Cloudformation template which is used to spin up resources in your account. Which of the following status in Cloudformation represents a failure.

A. UPDATE_COMPLETE_CLEANUPJN_PROGRESS
B. DELETE_COMPLETE
C. ROLLBACK_IN_PROGRESS
D. UPDATE_IN_PROGRESS

Answer: C

Explanation:
AWS Cloud Formation provisions and configures resources by making calls to the AWS services that are described in your template.

After all the resources have been created, AWS Cloud Formation reports that your stack has been created. You can then start using the resources in your stack. If stack creation fails, AWS CloudFormation rolls back your changes by deleting the resources that it created.
The below snapshot from Cloudformation shows what happens when there is an error in the stack creation.

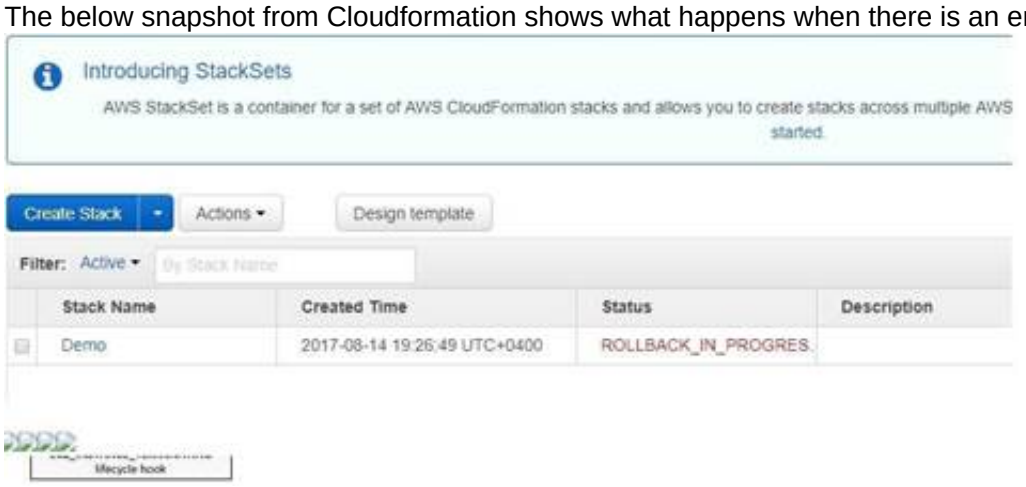

For more information on how Cloud Formation works, please refer to the below link: http://docs.ws.amazon.com/AWSCIoudFormation/latest/UserGuide/cfn-whatis-howdoesitwork-html

NEW QUESTION 79
You have a requirement to host a cluster of NoSQL databases. There is an expectation that there will be a lot of I/O on these databases. Which EBS volume type is best for high performance NoSQL cluster deployments?

A. io1
B. gp1
C. standard
D. gp2

Answer: A

Explanation:
Provisioned IOPS SSD should be used for critical business applications that require sustained IOPS performance, or more than 10,000 IOPS or 160 MiB/s of throughput per volume
This is ideal for Large database workloads, such as:
• MongoDB
• Cassandra
• MicrosoftSQL Server
• MySQL
• PostgreSQL
• Oracle
For more information on the various CBS Volume Types, please refer to the below link:
• http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/ CBSVolumeTvpes.html

NEW QUESTION 80
You run accounting software in the AWS cloud. This software needs to be online continuously during the day every day of the week, and has a very static requirement for compute resources. You also have other, unrelated batch jobs that need to run once per day at anytime of your choosing. How should you minimize cost?

A. Purchase a Heavy Utilization Reserved Instance to run the accounting softwar
B. Turn it off after hour
C. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.
D. Purch ase a Medium Utilization Reserved Instance to run the accounting softwar
E. Turn it off after hour
F. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.
G. Purchase a Light Utilization Reserved Instance to run the accounting softwar
H. Turn it off after hour
I. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.
J. Purch ase a Full Utilization Reserved Instance to run the accounting softwar
K. Turn it off after hour
L. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.

Answer: A

Explanation:
Reserved Instances provide you with a significant discount compared to On-Demand Instance pricing.
Reserved Instances are not physical instances, but rather a
billing discount applied to the use of On-Demand Instances in your account. These On-Demand Instances must match certain attributes in order to benefit from the billing discount
For more information, please refer to the below link:
• https://aws.amazon.com/about-aws/whats-new/2011/12/01/New-Amazon-CC2-Reserved- Instances-Options-Now-Available/
• https://aws.amazon.com/blogs/aws/reserved-instance-options-for-amazon-ec2/
• http://docs.aws.a mazon.com/AWSCC2/latest/UserGuide/ec2-reserved-instances.html Note:
It looks like these options are also no more available at present.
It looks like Convertible, Standard and scheduled are the new instance options. However the exams may still be referring to the old RIs.
https://aws.amazon.com/ec2/pricing/reserved-instances/

NEW QUESTION 83
You need to create a simple, holistic check for your system's general availablity and uptime. Your system presents itself as an HTTP-speaking API. What is the most simple tool on AWS to achieve this with?

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

Answer: A

Explanation:
Amazon Route 53 health checks monitor the health and performance of your web applications, web servers, and other resources. Each health check that you create
can monitor one of the following:
• The health of a specified resource, such as a web server
• The status of an Amazon Cloud Watch alarm
• The status of other health checks
For more information on Route53 Health checks, please refer to the below link:
• http://docs.aws.a mazon.com/Route53/latest/DeveloperGuide/dns-fa ilover.html

NEW QUESTION 86
You need to perform ad-hoc business analytics queries on well-structured data. Data comes in constantly at a high velocity. Your business intelligence team can understand SQL.
What AWS service(s) should you look to first?

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

Answer: B

Explanation:
Amazon Kinesis Firehose is the easiest way to load streaming data into AWS. It can capture, transform, and load streaming data into Amazon Kinesis Analytics, Amazon S3, Amazon Redshift, and Amazon Oasticsearch Sen/ice, enabling near real-time analytics with existing business intelligence tools and dashboards you're already using today. It is a fully managed service that automatically scales to match the throughput of your data and requires no ongoing administration. It can also batch, compress, and encrypt the data before loading it, minimizing the amount of storage used at the destination and increasing security.
For more information on Kinesis firehose, please visit the below URL:
• https://aws.amazon.com/kinesis/firehose/
Amazon Redshift is a fully managed, petabyte-scale data warehouse service in the cloud. You can start with just a few hundred gigabytes of data and scale to a petabyte or more. This enables you to use your data to acquire new insights for your business and customers. For more information on Redshift, please visit the below URL:
http://docs.aws.amazon.com/redshift/latest/mgmt/welcome.html

NEW QUESTION 87
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 91
You are planning on using encrypted snapshots in the design of your AWS Infrastructure. Which of the following statements are true with regards to EBS Encryption

A. Snapshottingan encrypted volume makes an encrypted snapshot; restoring an encrypted snapshot creates an encrypted volume when specified / requested.
B. Snapshotting an encrypted volume makes an encrypted snapshot when specified / requested; restoring an encrypted snapshot creates an encrypted volume when specified / requested.
C. Snapshotting an encrypted volume makes an encrypted snapshot; restoring an encrypted snapshot always creates an encrypted volume.
D. Snapshotting an encrypted volume makes an encrypted snapshot when specified / requested; restoring an encrypted snapshot always creates an encrypted volume.

Answer: C

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

• Data at rest inside the volume
• All data moving between the volume and the instance
• All snapshots created from the volume
Snapshots that are taken from encrypted volumes are automatically encrypted. Volumes that are created from encrypted snapshots are also automatically encrypted.
For more information on CBS encryption, please visit the below URL:
• http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/ CBSCncryption.html

NEW QUESTION 92
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 94
Which of the following tools does not directly support AWS OpsWorks, for monitoring your stacks?

A. AWSConfig
B. Amazon CloudWatch Metrics
C. AWSCIoudTrail
D. Amazon CloudWatch Logs

Answer: A

Explanation:
You can monitor your stacks in the following ways.
• AWS OpsWorks Stacks uses Amazon CloudWatch to provide thirteen custom metrics with detailed monitoring for each instance in the stack.
• AWS OpsWorks Stacks integrates with AWS CloudTrail to log every AWS OpsWorks Stacks API call and store the data in an Amazon S3 bucket.
• You can use Amazon CloudWatch Logs to monitor your stack's system, application, and custom logs.
For more information on Opswork monitoring, please visit the below URL:
• http://docs.aws.amazon.com/opsworks/latest/userguide/monitoring.html

NEW QUESTION 98
Your company wants to understand where cost is coming from in the company's production AWS account. There are a number of applications and services running at any given time. Without expending too much initial development time.how best can you give the business a good understanding of which applications cost the most per month to operate?

A. Create an automation script which periodically creates AWS Support tickets requesting detailed intra-month information about your bill.
B. Use custom CloudWatch Metrics in your system, and put a metric data point whenever cost is incurred.
C. Use AWS Cost Allocation Taggingfor all resources which support i
D. Use the Cost Explorer to analyze costs throughout the month.
E. Use the AWS Price API and constantly running resource inventory scripts to calculate total price based on multiplication of consumed resources over time.

Answer: C

Explanation:
A tag is a label that you or AWS assigns to an AWS resource. Each tag consists of a Areyand a value. A key can have more than one value. You can use tags to organize your resources, and cost allocation tags to track your AWS costs on a detailed level. After you activate cost allocation tags, AWS uses the cost allocation tags to organize your resource costs on your cost allocation report, to make it easier
for you to categorize and track your AWS costs. AWS provides two types of cost allocation tags, an A WS-generated tagand user-defined tags. AWS defines, creates, and applies the AWS-generated tag for you, and you define, create, and apply user-defined tags. You must activate both types of tags separately before they can appear in Cost Explorer or on a cost allocation report.
For more information on Cost Allocation tags, please visit the below URL: http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloctags.html

NEW QUESTION 99
You need your API backed by DynamoDB to stay online duringa total regional AWS failure. You can tolerate a couple minutes of lag or slowness during a large failure event, but the system should recover with normal operation after those few minutes. What is a good approach?

A. Set up DynamoDB cross-region replication in a master-standby configuration, with a single standby in another regio
B. Create an Auto Scaling Group behind an ELB in each of the two regions for your application layer in which DynamoDB is running i
C. Add a Route53 Latency DNS Record with DNS Failover, using the ELBs in the two regions as the resource records.
D. Set up a DynamoDB Global tabl
E. Create an Auto Scaling Group behind an ELB in each of the two regions for your application layer in which the DynamoDB is running i
F. Add a Route53 Latency DNS Record with DNS Failover, using the ELBs in the two regions as the resource records.

G. Set up a DynamoDB Multi-Region tabl
H. Create a cross-region ELB pointing to a cross-region Auto Scaling Group, and direct a Route53 Latency DNS Record with DNS Failover to the cross-region ELB.
I. Set up DynamoDB cross-region replication in a master-standby configuration, with a single standbyin another regio
J. Create a crossregion ELB pointing to a cross-region Auto Scaling Group, and direct a Route53 Latency DNS Record with DNS Failover to the cross- region ELB.

Answer: B

Explanation:
Updated based on latest AWS updates
Option A is invalid because using Latency based routing will sent traffic on the region with the standby instance. This is an active/passive replication and you can't write to the standby table unless there is a failover. Answer A can wort: only if you use a failover routing policy.
Option D is invalid because there is no concept of a cross region CLB.
Amazon DynamoDBglobal tables provide a fully managed solution for deploying a multi-region, multi-master database, without having to build and maintain your own replication solution. When you create a global table, you specify the AWS regions where you want the table to be available. DynamoDB performs all of the necessary tasks to create identical tables in these regions, and propagate ongoing data changes to all of them.
For more information on DynamoDB GlobalTables, please visit the below URL:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GlobalTables.html

NEW QUESTION 100
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 105
Your CTO has asked you to make sure that you know what all users of your AWS account are doing to change resources at all times. She wants a report of who is doing what over time, reported to her once per week, for as broad a resource type group as possible. How should you do this?

A. Create a global AWS CloudTrail Trai
B. Configure a script to aggregate the log data delivered to S3 once per week and deliver this to the CTO.
C. Use CloudWatch Events Rules with an SNS topic subscribed to all AWS API call
D. Subscribe the CTO to an email type delivery on this SNS Topic.
E. Use AWS 1AM credential reports to deliver a CSV of all uses of 1AM UserTokens overtime to the CTO.
F. Use AWS Config with an SNS subscription on a Lambda, and insert these changes over time into a DynamoDB tabl
G. Generate reports based on the contents of this table.

Answer: A

Explanation:
AWS CloudTrail is an AWS service that helps you enable governance, compliance, and operational and risk auditing of your AWS account. Actions taken by a user, role, or an AWS service are recorded as events in CloudTrail. Events include actions taken in the AWS Management Console, AWS Command Line Interface, and AWS SDKs and APIs.
Visibility into your AWS account activity is a key aspect of security and operational best practices. You can use CloudTrail to view, search, download, archive, analyze, and respond to account activity across your AWS infrastructure. You can identify who or what took which action, what resources were acted upon, when the event occurred, and other details to help you analyze and respond to activity in your AWS account.
For more information on Cloudtrail, please visit the below URL:
• http://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html

NEW QUESTION 110
You are building a mobile app for consumers to post cat pictures online. You will be storing the images in AWS S3. You want to run the system very cheaply and simply. Which one of these options allows you to build a photo sharing application with the right authentication/authorization implementation.

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

Answer: A

Explanation:
Amazon Cognito lets you easily add user sign-up and sign-in and manage permissions for your mobile and web apps. You can create your own user directory within Amazon Cognito. You can also choose to authenticate users through social identity providers such as Facebook, Twitter, or Amazon; with SAML identity solutions; or by using your own identity system. In addition, Amazon Cognito enables you to save data locally on users' devices, allowing your applications to work even when the devices are offline. You can then synchronize data across users' devices so that their app experience remains consistent regardless of the device

they use.
For more information on AWS Cognito, please visit the below URL:
• http://docs.aws.amazon.com/cognito/latest/developerguide/what-is-amazon-cognito.html

NEW QUESTION 112
You are building a Ruby on Rails application for internal, non-production use which uses MySQL as a database. You want developers without very much AWS experience to be able to deploy new code with a single command line push. You also want to set this up as simply as possible. Which tool is ideal for this setup?

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

Answer: D

Explanation:
With Elastic Beanstalk, you can quickly deploy and manage applications in the AWS Cloud without worrying about the infrastructure that runs those applications. AWS Elastic Beanstalk reduces management complexity without restricting choice or control. You simply upload your application, and Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring
Elastic Beanstalk supports applications developed in Java, PHP, .NET, Node.js, Python, and Ruby, as well as different container types for each language.
For more information on Elastic beanstalk, please visit the below URL:
• http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/Welcome.html

NEW QUESTION 116
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 118
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 122
You need to grant a vendor access to your AWS account. They need to be able to read protected messages in a private S3 bucket at their leisure. They also use AWS. What is the best way to accomplish this?

A. Create an 1AM User with API Access Key
B. Grant the User permissions to access the bucke
C. Give the vendor the AWS Access Key ID and AWS Secret Access Key for the User.
D. Create an EC2 Instance Profile on your accoun

E. Grant the associated 1AM role full access to the bucke
F. Start an EC2 instance with this Profile and give SSH access to the instance to the vendor.
G. Create a cross-account I AM Role with permission to access the bucket, and grant permission to use the Role to the vendor AWS account.D- Generate a signed S3 PUT URL and a signed S3 PUT URL, both with wildcard values and 2 year duration
H. Pass the URLs to the vendor.

Answer: C

Explanation:
You can use AWS Identity and Access Management (I AM) 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 Access, please visit the below URL:
• https://aws.amazon.com/blogs/security/tag/cross-account-access/

NEW QUESTION 127
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 132
Which of the following is the default deployment mechanism used by Elastic Beanstalk when the application is created via Console or EBCLI?

A. All at Once
B. Rolling Deployments
C. Rolling with additional batch
D. Immutable

Answer: B

Explanation:
The AWS documentation mentions
AWS Elastic Beanstalk provides several options for how deployments are processed, including deployment policies (All at once. Rolling, Rolling with additional batch,
and Immutable) and options that let you configure batch size and health check behavior during deployments. By default, your environment uses rolling deployments
if you created it with the console or EB CLI, or all at once deployments if you created it with a different client (API, SDK or AWS CLI).
For more information on Elastic Beanstalk deployments, please refer to the below link:
• http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.rolling-version- deploy.html

NEW QUESTION 137
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

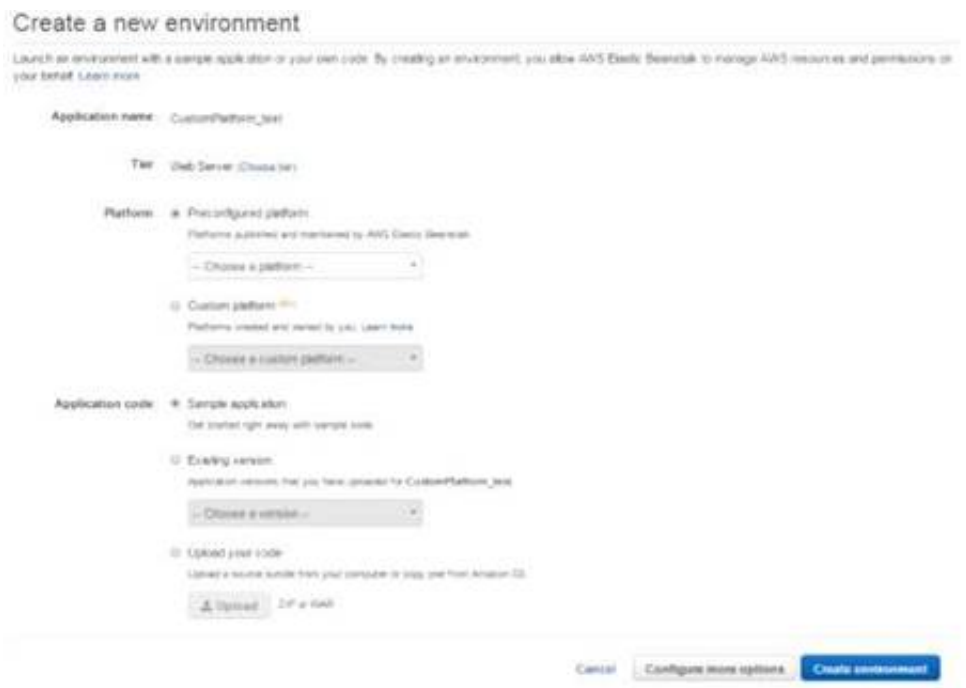

The high availability preset includes a load balancer; the low cost preset does not For more information on the configuration settings, please refer to the below link:
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environments-create-wizard.html

NEW QUESTION 138
You have an Autoscaling Group configured to launch EC2 Instances for your application. But you notice that the Autoscaling Group is not launching instances in the right proportion. In fact instances are being launched too fast. What can you do to mitigate this issue? Choose 2 answers from the options given below

A. Adjust the cooldown period set for the Autoscaling Group
B. Set a custom metric which monitors a key application functionality forthe scale-in and scale-out process.
C. Adjust the CPU threshold set for the Autoscaling scale-in and scale-out process.
D. Adjust the Memory threshold set forthe Autoscaling scale-in and scale-out process.

Answer: AB

Explanation:
The Auto Scaling cooldown period is a configurable setting for your Auto Scaling group that helps to ensure that Auto Scaling doesn't launch or terminate additional instances before the previous scaling activity takes effect.
For more information on the cool down period, please refer to the below link:
• http://docs^ws.a mazon.com/autoscaling/latest/userguide/Cooldown.html
Also it is better to monitor the application based on a key feature and then trigger the scale-in and scale-out feature accordingly. In the question, there is no mention of CPU or memory causing the issue.

NEW QUESTION 142
You have launched a cloudformation template, but are receiving a failure notification after the template was launched. What is the default behavior of Cloudformation in such a case

A. It will rollback all the resources that were created up to the failure point.
B. It will keep all the resources that were created up to the failure point.
C. It will prompt the user on whether to keep or terminate the already created resources
D. It will continue with the creation of the next resource in the stack

Answer: A

Explanation:
The AWS Documentation mentions
AWS Cloud Formation ensures all stack resources are created or deleted as appropriate. Because AWS CloudFormation treats the stack resources as a single unit,
they must all be created or deleted successfully for the stack to be created or deleted. If a resource cannot be created, AWS CloudFormation rolls the stack back and automatically deletes any resources that were created.
For more information on Cloudformation, please refer to the below link: http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/stacks.html

NEW QUESTION 143
When building a multicontainer Docker platform using Elastic Beanstalk, which of the following is required

A. DockerFile to create custom images during deployment
B. Prebuilt Images stored in a public or private online image repository.
C. Kurbernetes to manage the docker containers.
D. RedHatOpensift to manage the docker containers.

Answer: B

Explanation:
This is a special note given in the AWS Documentation for Multicontainer Docker platform for Elastic Beanstalk
Building custom images during deployment with a Dockerfile is not supported by the multicontainer Docker platform on Elastic Beanstalk. Build your images and deploy them to an online repository before creating an Elastic Beanstalk environment.
For more information on Multicontainer Docker platform for Elastic Beanstalk, please refer to the below link:
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_ecs.html

NEW QUESTION 147
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 148
Which of the following Deployment types are available in the CodeDeploy service. Choose 2 answers from the options given below

A. In-place deployment
B. Rolling deployment
C. Immutable deployment
D. Blue/green deployment

Answer: AD

Explanation:
The following deployment types are available
1. In-place deployment: The application on each instance in the deployment group is stopped, the latest application revision is installed, and the new version of the application is started and validated.
2. Blue/green deployment: The instances in a deployment group (the original environment) are replaced by a different set of instances (the replacement environment)
For more information on Code Deploy please refer to the below link:
• http://docs.aws.amazon.com/codedeploy/latest/userguide/primary-components.html

NEW QUESTION 153
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 155
Which of the following Cache Engines does Opswork have built in support for?

A. Redis
B. Memcache
C. Both Redis and Memcache
D. There is no built in support as of yet for any cache engine

Answer: B

Explanation:
The AWS Documentation mentions
AWS OpsWorks Stacks provides built-in support for Memcached. However, if Redis better suits your requirements, you can customize your stack so that your application servers use OastiCache Redis. Although it works with Redis clusters, AWS clearly specifies that AWS Opsworks stacks provide built in support for Memcached.
Amazon OastiCache is an AWS service that makes it easy to provide caching support for your application server, using either the Memcached or Redis caching engines. OastiCache can be used to improve the application server performance running on AWS Opsworks stacks.
For more information on Opswork and Cache engines please refer to the below link:
• http://docs^ws.a mazon.com/opsworks/latest/userguide/other-se rvices-redis.htm l

NEW QUESTION 159
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 164
You have ana video processing application hosted in AWS. The video's are uploaded by users onto the site. You have a program that is custom built to process those videos. The program is able to recover incase there are any failures when processing the videos. Which of the following mechanisms can be used to deploy the instances for carrying out the video processing activities, ensuring that the cost is kept at a minimum.

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

Answer: B

Explanation:
Since the application can recover from failures and cost is the priority, then Spot instances are the best bet for this requirement. The launch configuration has the facility to request for Spot Instances.
The below snapshot from the Launch configuration section shows that Spot Instances can be used for AutoScaling Groups.

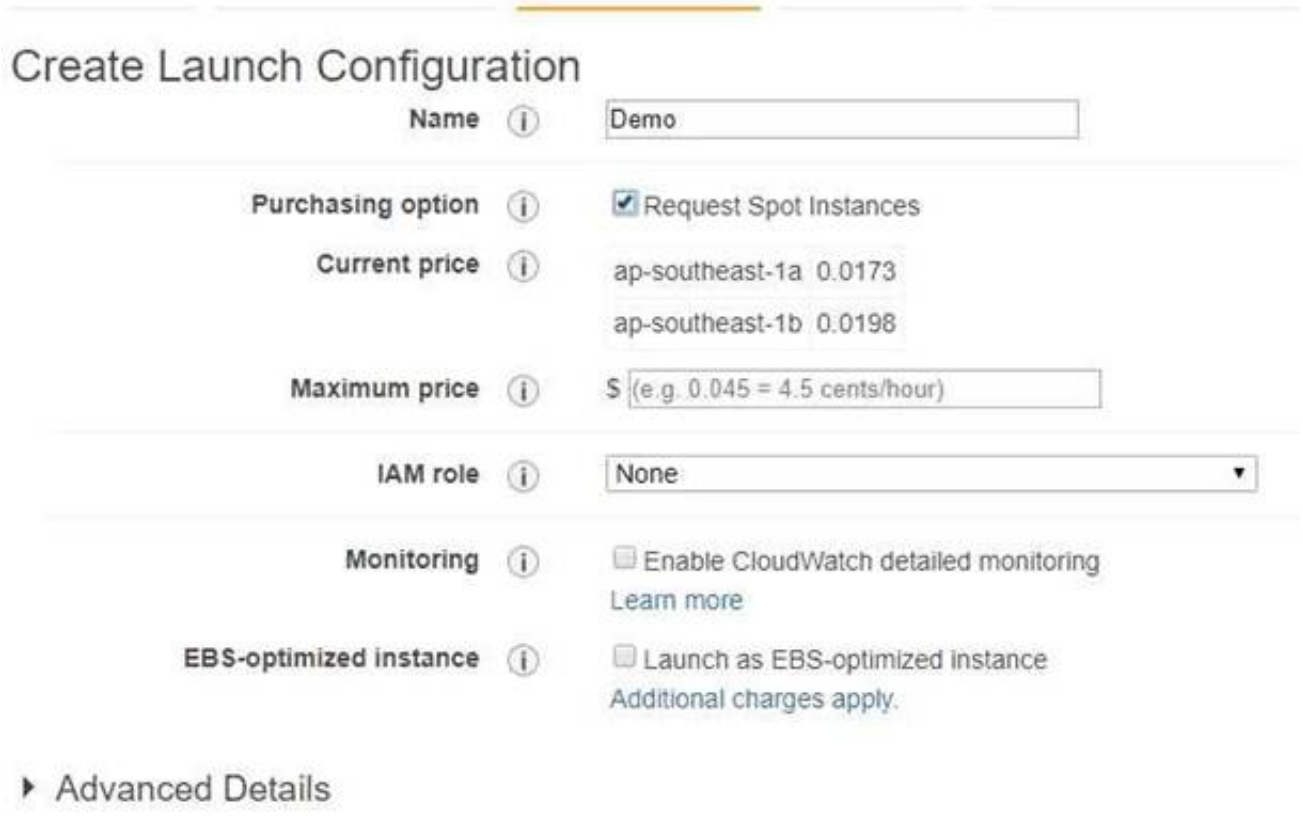

For more information on Spot Instances and Autoscaling, please visit the below URL:
• http://docs^ws.amazon.com/autoscaling/latest/userguide/US-SpotInstances.html

NEW QUESTION 168
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 173
Your company has the requirement to set up instances running as part of an Autoscaling Group. Part of the requirement is to use Lifecycle hooks to setup custom based software's and do the necessary configuration on the instances. The time required for this setup might take an hour, or might finish before the hour is up. How should you setup lifecycle hooks for the Autoscaling Group. Choose 2 ideal actions you would include as part of the lifecycle hook.

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

Answer: AD

Explanation:
The AWS Documentation provides the following information on lifecycle hooks
By default, the instance remains in a wait state for one hour, and then Auto Scaling continues the launch or terminate process (Pending: Proceed or Terminating: Proceed). If you need more time, you can restart the timeout period by recording a heartbeat. If you finish before the timeout period ends, you can complete the lifecycle action, which continues the launch or termination process
For more information on AWS Lifecycle hooks, please visit the below URL:
• http://docs.aws.amazon.com/autoscaling/latest/userguide/lifecycle-hooks.html

NEW QUESTION 175
Your application is having a very high traffic, so you have enabled autoscaling in multi availability zone to suffice the needs of your application but you observe that one of the availability zone is not receiving any traffic. What can be wrong here?

A. Autoscaling only works for single availability zone
B. Autoscaling can be enabled for multi AZ only in north Virginia region
C. Availability zone is not added to Elastic load balancer
D. Instances need to manually added to availability zone

Answer: C

Explanation:
When you add an Availability Zone to your load balancer. Elastic Load Balancing creates a load balancer node in the Availability Zone. Load balancer nodes accept traffic from clients and forward requests to the healthy registered instances in one or more Availability Zones.
For more information on adding AZ's to CLB, please refer to the below U RL:
http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/enable-disable-az.htmI

NEW QUESTION 179
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 183
You have a legacy application running that uses an m4.large instance size and cannot scale with Auto Scaling, but only has peak performance 5% of the time. This is a huge waste of resources and money so your Senior Technical Manager has set you the task of trying to reduce costs while still keeping the legacy application running as it should. Which of the following would best accomplish the task your manager has set you? Choose the correct answer from the options below

A. Use a T2burstable performance instance.

B. Use a C4.large instance with enhanced networking.
C. Use two t2.nano instances that have single Root I/O Visualization.
D. Use t2.nano instance and add spot instances when they are required.

Answer: A

Explanation:
The aws documentation clearly indicates using T2 CC2 instance types for those instances which don't use CPU that often.
T2
T2 instances are Burstable Performance Instances that provide a baseline level of CPU performance with the ability to burst above the baseline.
T2 Unlimited instances can sustain high CPU performance for as long as a workload needs it. For most general-purpose workloads, T2 Unlimited instances will provide ample performance without any additional charges. If the instance needs to run at higher CPU utilization for a prolonged period, it can also do so at a flat additional charge of 5 cents per vCPU-hour.
The baseline performance and ability to burst are governed by CPU Credits. T2 instances receive CPU Credits continuously at a set rate depending on the instance size, accumulating CPU Credits when they are idle, and consuming CPU credits when they are active. T2 instances are a good choice for a variety of general-purpose workloads including micro-services, low-latency interactive applications, small and medium databases, virtual desktops, development, build and stage environments, code repositories, and product prototypes. For more information see Burstable Performance Instances.
For more information on F_C2 instance types please see the below link: https://aws.amazon.com/ec2/instance-types/

NEW QUESTION 184
You're building a mobile application game. The application needs permissions for each user to communicate and store data in DynamoDB tables. What is the best method for granting each mobile device that installs your application to access DynamoDB tables for storage when required? Choose the correct answer from the options below

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

Answer: C

Explanation:
Answer - C
For access to any AWS service, the ideal approach for any application is to use Roles. This is the first preference.
For more information on 1AM policies please refer to the below link:
http://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html
Next for any web application, you need to use web identity federation. Hence option D is the right option. This along with the usage of roles is highly stressed in the aws documentation.
The AWS documentation mentions the following
When developing a web application it is recommend not to embed or distribute long-term AWS credentials with apps that a user downloads to a device, even in an encrypted store. Instead, build your app so that it requests temporary AWS security credentials dynamically when needed using web identity federation. The supplied temporary credentials map to an AWS role that has only the permissions needed to perform the tasks required by the mobile app.
For more information on web identity federation please refer to the below link: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html

NEW QUESTION 188
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

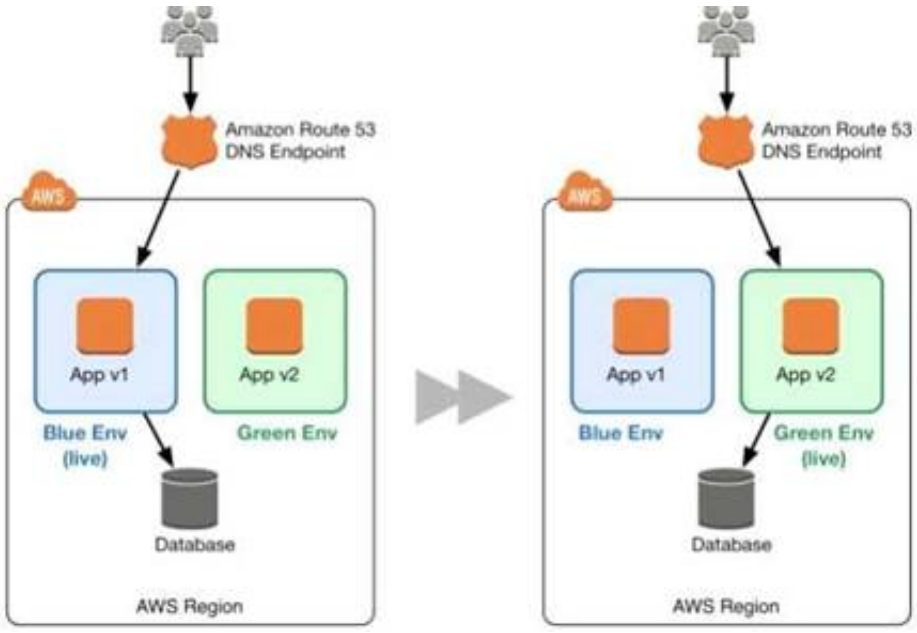

Please find the below link on a white paper for blue green deployments
• https://d03wsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 189
You are working as an AWS Devops admins for your company. You are in-charge of building the infrastructure for the company's development teams using Cloudformation. The template will include building the VPC and networking components, installing a LAMP stack and securing the created resources. As per the AWS best practices what is the best way to design this template

A. Create a single cloudformation template to create all the resources since it would be easierfrom the maintenance perspective.
B. Create multiple cloudformation templates based on the number of VPC's in the environment.
C. Create multiple cloudformation templates based on the number of development groups in the environment.
D. Create multiple cloudformation templates for each set of logical resources, one for networking, the otherfor LAMP stack creation.

Answer: D

Explanation:
Creating multiple cloudformation templates is an example of using nested stacks. The advantage of using nested stacks is given below as per the AWS documentation
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 practices, please refer to the below link: http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/best-practices.html

NEW QUESTION 194
Your security officer has told you that you need to tighten up the logging of all events that occur on your AWS account. He wants to be able to access all events that occur on the account across all regions quickly and in the simplest way possible. He also wants to make sure he is the only person that has access to these events in the most secure way possible. Which of the following would be the best solution to assure his requirements are met? Choose the correct answer from the options below

A. Use CloudTrail to logall events to one S3 bucke
B. Make this S3 bucket only accessible by your security officer with a bucket policy that restricts access to his user only and also add MFA to the policy for a further level of securit
C. ^/
D. Use CloudTrail to log all events to an Amazon Glacier Vaul
E. Make sure the vault access policy only grants access to the security officer's IP address.
F. Use CloudTrail to send all API calls to CloudWatch and send an email to the security officer every time an API call is mad
G. Make sure the emails are encrypted.
H. Use CloudTrail to log all events to a separate S3 bucket in each region as CloudTrail cannot write to a bucket in a different regio
I. Use MFA and bucket policies on all the different buckets.

Answer: A

Explanation:
AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account. With CloudTrail, you can log, continuously monitor, and retain events related to API calls across your AWS infrastructure. CloudTrail provides a history of AWS API calls for your account, including API calls made through the AWS Management Console, AWS SDKs, command line tools, and other AWS services. This history simplifies security analysis, resource change tracking, and troubleshooting.
You can design cloudtrail to send all logs to a central S3 bucket. For more information on cloudtrail, please visit the below URL:
? https://aws.amazon.com/cloudtrail/

NEW QUESTION 195
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 197
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 202
You are having a web and worker role infrastructure defined in AWS using Amazon EC2 resources. You are using SQS to manage the jobs being send by the web role. Which of the following is the right way to ensure the worker processes are adequately setup to handle the number of jobs send by the web role

A. Use Cloudwatch monitoring to check the size of the queue and then scale out SQS to ensure that it can handle the right number of jobs
B. Use ELB to ensure that the load is evenly distributed to the set of web and worker instances
C. Use Route53 to ensure that the load is evenly distributed to the set of web and worker instances
D. Use Cloudwatch monitoring to check the size of the queue and then scale out using Autoscaling to ensure that it can handle the right number of jobs

Answer: D

Explanation:
The below diagram shows how SGS can be used to manage the communication between the Web and worker roles. The number of messages in the SQS queue can be used to determine the number of instances that should be there in the AutoScaling Group.

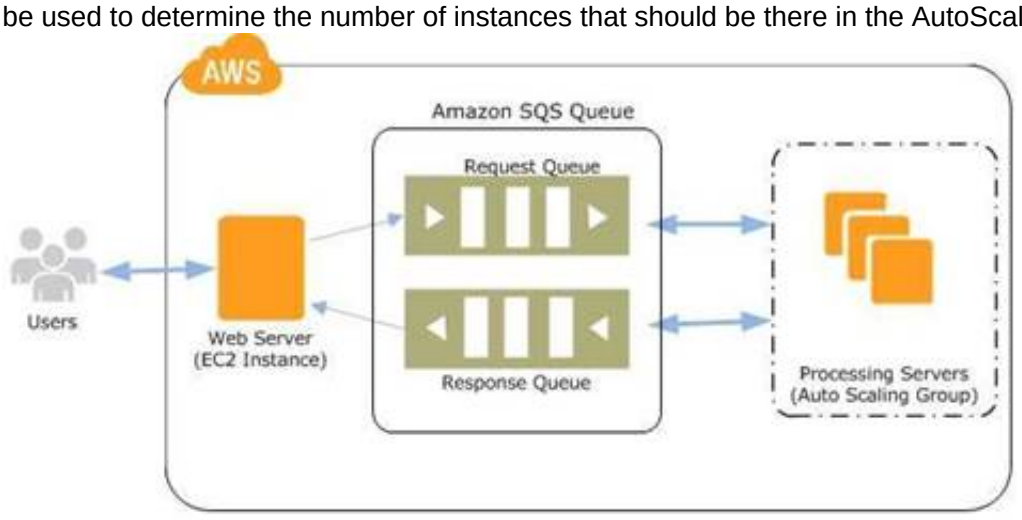

For more information on SQS and Autoscaling, please refer to the below U RL: http://docs.aws.amazon.com/autoscaling/latest/userguide/as-using-sqs-queue.html

NEW QUESTION 203
Which of the following is incorrect when it comes to using the instances in an Opswork stack?

A. In a stack you can use a mix of both Windowsand Linux operating systems
B. You can start and stop instances manually in a stack
C. You can use custom AMI'S as long as they are based on one of the AWS OpsWorks Stacks- supported AMIs
D. You can use time-based automatic scaling with any stack

Answer: A

Explanation:
The AWS documentation mentions the following about Opswork stack
• 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.
• You can use custom AMIs (Amazon Machine Images), but they must be based on one of the AWS Ops Works Stacks-supported AMIs
• You can start and stop instances manually or have AWS OpsWorks 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.
• In addition to using AWS OpsWorks Stacks to create Amazon EC2 instances, you can also register instances with a Linux stack that were created outside of AWS OpsWorks Stacks.
For more information on Opswork stacks, please visit the below link: http://docs.aws.amazon.com/opsworks/latest/userguide/workinginstances-os.html

NEW QUESTION 204
Which of the following will you need to consider so you can set up a solution that incorporates single sign-on from your corporate AD or LDAP directory and restricts access for each user to a designated user folder in a bucket? Choose 3 Answers from the options below

A. Setting up a federation proxy or identity provider
B. Using AWS Security Token Service to generate temporary tokens
C. Tagging each folder in the bucket
D. Configuring 1AM role
E. Setting up a matching 1AM user for every user in your corporate directory that needs access to a folder in the bucket

Answer: ABD

Explanation:
The below diagram showcases how authentication is carried out when having an identity broker. This is an example of a SAML connection, but the same concept holds true for getting access to an AWS resource.

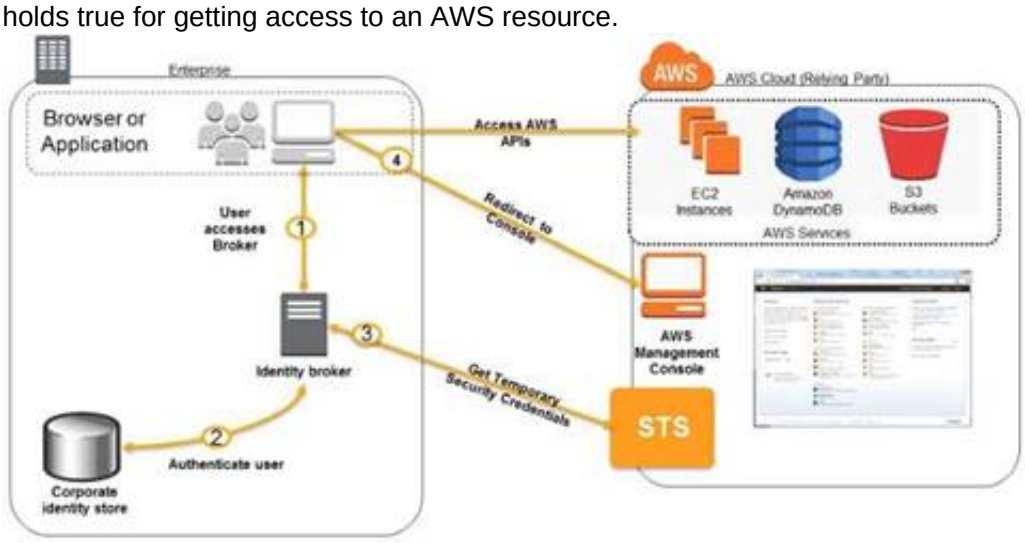

For more information on federated access, please visit the below link: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_com mon-scenarios_federated-users.htm I
https://docs.aws.a mazon.com/IAM/latest/UserGuide/id_roles_create_for-idp_sam I. html?icmpid=docs_iam_console
https://aws.ama zon.com/blogs/secu rity/writing-iam-policies-grant-access-to-user-specific-folders-in-an-amazon-s3-bucket/

NEW QUESTION 207
A custom script needs to be passed to a new Amazon Linux instances created in your Auto Scalinggroup. Which feature allows you to accomplish this?

A. User data
B. EC2Config service
C. 1AM roles
D. AWSConfig

Answer: A

Explanation:
When you configure an instance during creation, you can add custom scripts to the User data section. So in Step 3 of creating an instance, in the Advanced Details section, we can enter custom scripts in the User Data section. The below script installs Perl during the instance creation of the CC2 instance.

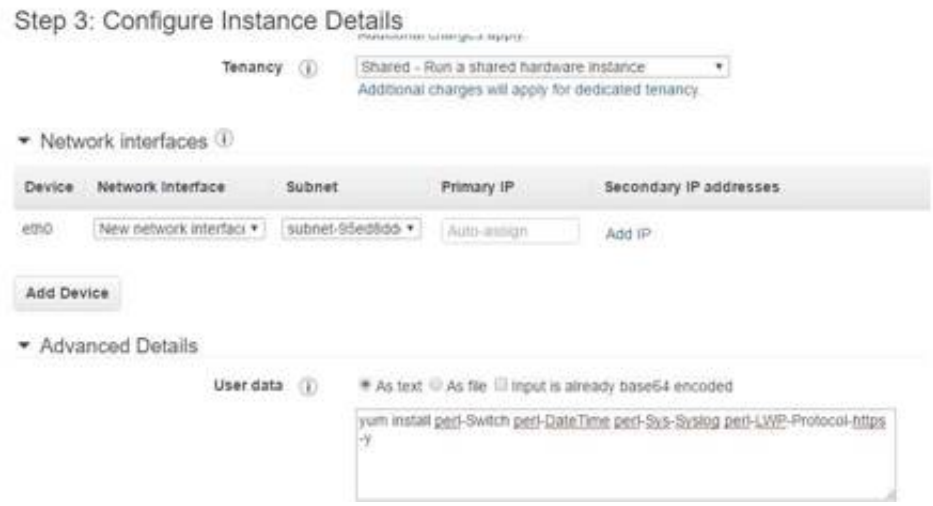

For more information on user data please refer to the URL:
• http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/ec2-i nstance-metadata.htm I

NEW QUESTION 210
You are planning on using AWS Code Deploy in your AWS environment. Which of the below features of AWS Code Deploy can be used to Specify scripts to be run on each instance at various stages of the deployment process

A. AppSpecfile
B. CodeDeployfile
C. Configfile
D. Deploy file

Answer: A

Explanation:
The AWS Documentation mentions the following on AWS Code Deploy
An application specification file (AppSpec file), which is unique to AWS CodeDeploy, is a YAML- formatted file used to:
Map the source files in your application revision to their destinations on the instance. Specify custom permissions for deployed files.
Specify scripts to be run on each instance at various stages of the deployment process. For more information on AWS CodeDeploy, please refer to the URL:
http://docs.aws.amazon.com/codedeploy/latest/userguide/application-specification-files.html

NEW QUESTION 214
A user is trying to save some cost on the AWS services. Which of the below mentioned options will not help him save cost?

A. Delete the unutilized EBS volumes once the instance is terminated
B. Delete the AutoScaling launch configuration after the instances are terminated
C. Release the elastic IP if not required once the instance is terminated
D. Delete the AWS ELB after the instances are terminated

Answer: B

Explanation:
Option A is wrong because CBS volumes does have a costing aspect and hence deleting the volumes will save on cost
Option C is wrong because Elastic IP will consume cost if not removed. Option D is wrong because CLB also incur costs.
Only Autoscaling groups are free of cost. It's only the underlying resources which you are charged for. For more information on AWS Pricing, please visit the link:
https://aws.amazon.com/pricing/services/

NEW QUESTION 215
An organization is planning to use AWS for their Production Rollout. The organizations wants to implement automation for deployment, such that it will automatically create a LAMP stack, deploy an RDS MySQLDB instance, download the latest PHP installable from S3 and set up the ELB. Which of the below mentioned AWS services meets the requirement for making an orderly deployment of the software?

A. AWS Elastic Beanstalk
B. AWSCloudfront
C. AWS Cloudformation
D. AWS DevOps

Answer: C

Explanation:
When you want to automate deployment, the automatic choice is Cloudformation. Below is the excerpt from AWS on cloudformation.
AWS Cloud Formation gives developers and systems administrators an easy way to create and manage a collection of related AWS resources, provisioning updating them in an orderly and predictable fashion.
You can use AWS Cloud Formation's sample templates or create your own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run your application. You don't need to figure out the order for provisioning AWS services or the subtleties of making those dependencies work. Cloud Formation takes care of this for you. After the AWS resources are deployed, you can modify and update them in a controlled and predictable way, in effect applying version control to your AWS infrastructure the same way you do with your software For more information on Cloud Formation, please visit the link:
• https://aws.amazon.com/cloudformation/ As per AWS,
"AWS Clastic Beanstalk provides support for running Amazon Relational Database Service (Amazon RDS) instances in your Clastic Beanstalk environment. This works great for development and testing environments. However, it isn't ideal for a production environment because it ties the lifecycle of the database instance to the lifecycleofyour application's environment."
• https://docs.aws.a mazon.com/elasticbeanstalk/latesWdg/AW SHowTo.RDS.htm I

NEW QUESTION 220
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 222
.....

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