Exam Questions DOP-C01

AWS Certified DevOps Engineer- Professional

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

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

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

Answer: B

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

NEW QUESTION 2
You currently have the following setup in AWS
1) An Elastic Load Balancer
2) Auto Scaling Group which launches EC2 Instances
3) AMIs with your code pre-installed
You want to deploy the updates of your app to only a certain number of users. You want to have a cost-effective solution. You should also be able to revert back quickly. Which of the below solutions is the most feasible one?

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

Answer: A

Explanation:
The Weighted Routing policy of Route53 can be used to direct a proportion of traffic to your application. The best option is to create a second CLB, attach the new Autoscaling Group and then use Route53 to divert the traffic.
Option B is wrong because just having EC2 instances running with the new code will not help.
Option C is wrong because Clastic beanstalk is good for development environments, and also there is no mention of having 2 environments where environment url's
can be swapped.
Option D is wrong because you still need Route53 to split the traffic.
For more information on Route53 routing policies, please refer to the below link: http://docs.aws.a mazon.com/Route53/latest/DeveloperGuide/routing-policy. html

NEW QUESTION 3
Your application is currently running on Amazon EC2 instances behind a load balancer. Your management has decided to use a Blue/Green deployment strategy. How should you implement this for each deployment?

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

Answer: C

Explanation:
The below diagram shows how this can be done

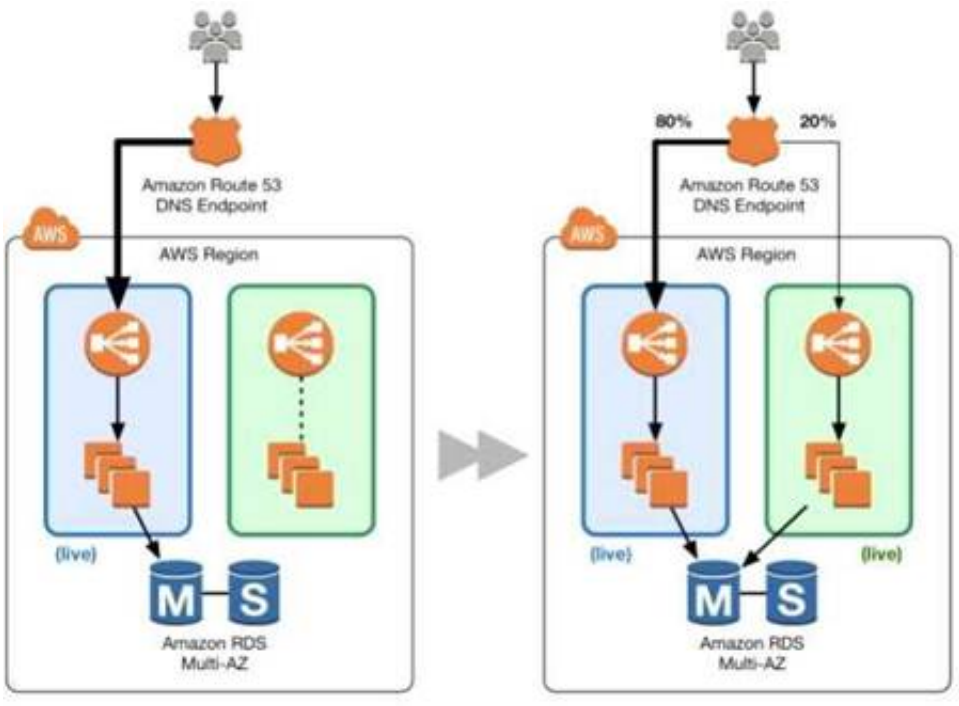

1) First create a new ELB which will be used to point to the new production changes.
2) Use the Weighted Route policy for Route53 to distribute the traffic to the 2 ELB's based on a 80- 20% traffic scenario. This is the normal case, the % can be changed based on the requirement.
3) Finally when all changes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

NEW QUESTION 4
Which Auto Scaling process would be helpful when testing new instances before sending traffic to them, while still keeping them in your Auto Scaling Group?

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

Answer: D

Explanation:
If you suspend Ad dTo Load Balancer, Auto Scaling launches the instances but does not add them to the load balancer or target group. If you resume the AddTo Load Balancer process. Auto Scaling resumes adding instances to the load balancer or target group when they are launched. However, Auto Scaling does
not add the instances that were launched while this process was suspended. You must register those
instances manually.
Option A is invalid because this just balances the number of CC2 instances in the group across the Availability Zones in the region
Option B is invalid because this just checks the health of the instances. Auto Scaling marks an instance as unhealthy if Amazon CC2 or Clastic Load Balancing tells
Auto Scaling that the instance is unhealthy.
Option C is invalid because this process just terminates instances that are marked as unhealthy and later creates new instances to replace them.
For more information on process suspension, please refer to the below document link: from AWS http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resume-processes.html

NEW QUESTION 5
Your company has multiple applications running on AWS. Your company wants to develop a tool that notifies on-call teams immediately via email when an alarm is triggered in your environment. You have multiple on-call teams that work different shifts, and the tool should handle notifying the correct teams at the correct times. How should you implement this solution?

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

Answer: D

Explanation:
Option D fulfils all the requirements

1) First is to create a SNS topic for each group so that the required members get the email addresses.
2) Ensure the application uses the HTTPS endpoint and the SDK to publish messages Option A is invalid because the SQS service is not required.
Option B and C are incorrect. As per the requirement we need to provide notification to only those on-call teams who are working in that particular shift when an alarm is triggered. It need not have to be send to all the on-call teams of the company. With Option B & C, since we are not configuring the SNS topic for each on call team the notifications will be send to all the on-call teams. Hence these 2 options are invalid. For more information on setting up notifications, please refer to the below document link: from AWS http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US_SetupSNS.html

NEW QUESTION 6
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/WhatIsCloudWatchLogs.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
Your mobile application includes a photo-sharing service that is expecting tens of thousands of users at launch. You will leverage Amazon Simple Storage Service (S3) for storage of the user Images, and you must decide how to authenticate and authorize your users for access to these images. You also need to manage the storage of these images. Which two of the following approaches should you use? Choose two answers from the options below

A. Create an Amazon S3 bucket per user, and use your application to generate the S3 URI for the appropriate content.
B. Use AWS Identity and Access Management (1AM) user accounts as your application-level user database, and offload the burden of authentication from your application code.
C. Authenticate your users at the application level, and use AWS Security Token Service (STS) to grant token-based authorization to S3 objects.
D. Authenticate your users at the application level, and send an SMS token message to the use
E. Create an Amazon S3 bucket with the same name as the SMS message token, and move the user's objects to that bucket.
F. Use a key-based naming scheme comprised from the user IDs for all user objects in a single Amazon S3 bucket.

Answer: CE

Explanation:
The AWS Security Token Service (STS) is a web service that enables you to request temporary, limited-privilege credentials for AWS Identity and Access Management (1AM) users or for users that you authenticate (federated users). The token can then be used to grant access to the objects in S3.
You can then provides access to the objects based on the key values generated via the user id. Option A is possible but then becomes a maintenance overhead because of the number of buckets. Option B is invalid because 1AM users is not a good security practice.
Option D is invalid because SMS tokens are not efficient for this requirement.
For more information on the Security Token Service please refer to the below link: http://docs.aws.amazon.com/STS/latest/APIReference/Welcome.htrril

NEW QUESTION 8
You have an Auto Scaling group with 2 AZs. One AZ has 4 EC2 instances and the other has 3 EC2 instances. None of the instances are protected from scale in. Based on the default Auto Scaling termination policy what will happen?

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

Answer: D

Explanation:
The default termination policy is designed to help ensure that your network architecture spans Availability Zones evenly. When using the default termination policy. Auto Scaling selects an instance to terminate as follows:
Auto Scaling determines whether there are instances in multiple Availability Zones. If so, it selects the Availability Zone with the most instances and at least one instance that is not protected from scale in. If there is more than one Availability Zone with this number of instances. Auto Scaling selects the Availability Zone with the instances that use the oldest launch configuration. For more information on Autoscaling instance termination please refer to the below link: http://docs.aws.amazon.com/autoscaling/latest/userguide/as-instance-termination.html

NEW QUESTION 9
You are administering a continuous integration application that polls version control for changes and then launches new Amazon EC2 instances for a full suite of build tests. What should you do to ensure the lowest overall cost while being able to run as many tests in parallel as possible?

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

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

Answer: B

Explanation:
Continuous Integration (CI) is a development practice that requires developers to integrate code into a shared repository several times a day. Each check-in is then verified by an automated build, allowing teams to detect problems early.
Option A and D are invalid because you can do build tests on a CI system and not only Syntax tests. And Syntax tests are normally done during coding time and not during the build time.
Option C is invalid because Opswork is ideally not used for build and integration tests.
For an example of a Continuous integration system, please refer to the Jenkins system via the url below
• https://jenkins.io/

NEW QUESTION 10
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 10
You have an Auto Scaling group with an Elastic Load Balancer. You decide to suspend the Auto Scaling AddToLoadBalancer for a short period of time. What will happen to the instances launched during the suspension period?

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

Answer: C

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

NEW QUESTION 11
You have a current Clouformation template defines in AWS. You need to change the current alarm threshold defined in the Cloudwatch alarm. How can you achieve this?

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

Answer: C

Explanation:
Option A is incorrect because Cloudformation templates have the option to update resources.
Option B is incorrect because only those resources that need to be changed as part of the stack update are actually updated.
Option D is incorrect because deleting the stack is not the ideal option when you already have a change option available.
When you need to make changes to a stack's settings or change its resources, you update the stack instead of deleting it and creating a new stack. For example, if you
have a stack with an EC2 instance, you can update the stack to change the instance's AMI ID.
When you update a stack, you submit changes, such as new input parameter values or an updated template. AWS CloudFormation compares the changes you submit with the current state of your stack and updates only the changed resources
For more information on stack updates please refer to the below link:
• http://docs.aws.a mazon.com/AWSCIoudFormation/latest/UserGuide/using-cfn-updating- stacks.htmI

NEW QUESTION 15

You are using CloudFormation to launch an EC2 instance and then configure an application after the instance is launched. You need the stack creation of the ELB and Auto Scaling to wait until the EC2 instance is launched and configured properly. How do you do this?

A. It is not possible for the stack creation to wait until one service is created and launched
B. Use the WaitCondition resource to hold the creation of the other dependent resources
C. Use a CreationPolicy to wait for the creation of the other dependent resources >/
D. Use the HoldCondition resource to hold the creation of the other dependent resources

Answer: C

Explanation:
When you provision an Amazon EC2 instance in an AWS Cloud Formation stack, you might specify additional actions to configure the instance, such as install software packages or bootstrap applications. Normally, CloudFormation proceeds with stack creation after the instance has been successfully created. However, you can use a Creation Pol icy so that CloudFormation proceeds with stack creation only after your configuration actions are done. That way you'll know your applications are ready to go after stack creation succeeds.
A Creation Policy instructs CloudFormation to wait on an instance until CloudFormation receives the specified number of signals
Option A is invalid because this is possible
Option B is invalid because this is used make AWS CloudFormation pause the creation of a stack and wait for a signal before it continues to create the stack
For more information on this, please visit the below URL:
• https://aws.amazon.com/blogs/devops/use-a-creationpolicy-to-wait-for-on-instance- configurations/

NEW QUESTION 20
One of the instances in your Auto Scaling group health check returns the status of Impaired to Auto Scaling. What will Auto Scaling do in this case.

A. Terminate the instance and launch a new instance
B. Send an SNS notification
C. Perform a health check until cool down before declaring that the instance has failed
D. Wait for the instance to become healthy before sending traffic

Answer: A

Explanation:
Auto Scaling periodically performs health checks on the instances in your Auto Scaling group and identifies any instances that are unhealthy. You can configure Auto Scaling to determine the health status of an instance using Amazon EC2 status checks. Clastic Load Balancing health checks, or custom health checks
By default. Auto Scaling health checks use the results of the CC2 status checks to determine the health status of an instance. Auto Scaling marks an instance as unhealthy if its instance fails one or more of the status checks.
For more information monitoring in Autoscaling, please visit the below URL: http://docs.aws.a mazon.com/autoscaling/latest/userguide/as-mon itoring-features.html

NEW QUESTION 24
After a daily scrum with your development teams, you've agreed that using Blue/Green style deployments would benefit the team. Which technique should you use to deliver this new requirement?

A. Re-deploy your application on AWS Elastic Beanstalk, and take advantage of Elastic Beanstalk deployment types.
B. Using an AWS CloudFormation template, re-deploy your application behind a load balancer, launch a new AWS CloudFormation stack during each deployment, update your load balancer to send half your traffic to the new stack while you test, after verification update the load balancer to send 100% of traffic to the new stack, and then terminate the old stack.
C. Create a new Autoscaling group with the new launch configuration and desired capacity same as that of the initial Autoscaling group andassociate it with the same load balance
D. Once the new AutoScaling group's instances got registered with ELB, modify the desired capacity of the initial AutoScal ing group to zero and gradually delete the old Auto scaling grou
E. •>/
F. Using an AWS OpsWorks stack, re-deploy your application behind an Elastic Load Balancing load balancer and take advantage of OpsWorks stack versioning, during deployment create a new version of your application, tell OpsWorks to launch the new version behind your load balancer, and when the new version is launched, terminate the old OpsWorks stack.

Answer: C

Explanation:
This is given as a practice in the Green Deployment Guides

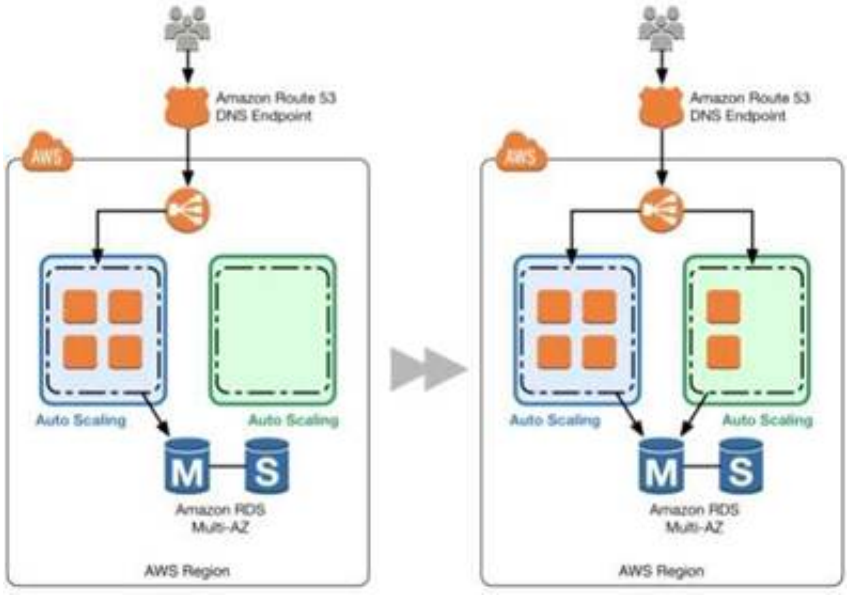

A blue group carries the production load while a green group is staged and deployed with the new code. When if s time to deploy, you simply attach the green group to
the existing load balancer to introduce traffic to the new environment. For HTTP/HTTP'S listeners, the load balancer favors the green Auto Scaling group because it uses a least outstanding requests routing algorithm

As you scale up the green Auto Scaling group, you can take blue Auto Scaling group instances out of service by either terminating them or putting them in Standby state.
For more information on Blue Green Deployments, please refer to the below document link: from AWS
https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 25
You have been tasked with deploying a scalable distributed system using AWS OpsWorks. Your distributed system is required to scale on demand. As it is distributed, each node must hold a configuration file that includes the hostnames of the other instances within the layer. How should you configure AWS OpsWorks to manage scaling this application dynamically?

A. Create a Chef Recipe to update this configuration file, configure your AWS OpsWorks stack to use custom cookbooks, and assign this recipe to the Configure LifeCycle Event of the specific layer.
B. Update this configuration file by writing a script to poll the AWS OpsWorks service API for new instance
C. Configure your base AMI to execute this script on Operating System startup.
D. Create a Chef Recipe to update this configuration file, configure your AWS OpsWorks stack to use custom cookbooks, and assign this recipe to execute when instances are launched.
E. Configure your AWS OpsWorks layer to use the AWS-provided recipe for distributed host configuration, and configure the instance hostname and file path parameters in your recipes settings.

Answer: A

Explanation:
Please check the following AWS DOCs which provides details on the scenario. Check the example of "configure".
? https://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-events.html You can use the Configure Lifecycle event
This event occurs on all of the stack's instances when one of the following occurs:
• An instance enters or leaves the online state.
• You associate an Elastic IP address with an instance or disassociate one from an instance.
• You attach an Elastic Load Balancing load balancer to a layer, or detach one from a layer. Ensure the Opswork layer uses a custom Cookbook

For more information on Opswork stacks, please refer to the below document link: from AWS
• http://docs.aws.a mazon.com/opsworks/latest/userguide/welcome_classic.html

NEW QUESTION 29
You have a large number of web servers in an Auto Scalinggroup behind a load balancer. On an hourly basis, you want to filter and process the logs to collect data on unique visitors, and then put that data in a durable data store in order to run reports. Web servers in the Auto Scalinggroup are constantly launching and terminating based on your scaling policies, but you do not want to lose any of the log data from these servers during a stop/termination initiated by a user or by Auto Scaling. What two approaches will meet these requirements? Choose two answers from the optionsgiven below.

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

Answer: AC

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

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

NEW QUESTION 32
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 35
Your company develops a variety of web applications using many platforms and programming languages with different application dependencies. Each application must be developed and deployed quickly and be highly available to satisfy your business requirements. Which of the following methods should you use to deploy these applications rapidly?

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

Answer: A

Explanation:
Elastic Beanstalk supports the deployment of web applications from Docker containers. With Docker containers, you can define your own runtime environment. You can choose your own platform, programming language, and any application dependencies (such as package managers or tools), that aren't supported by other platforms. Docker containers are self-contained and include all the configuration information and software your web application requires to run.
By using Docker with Elastic Beanstalk, you have an infrastructure that automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring.
For more information on Dockers and Elastic beanstalk please refer to the below link:
• http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker.html

NEW QUESTION 39
As an architect you have decided to use CloudFormation instead of OpsWorks or Elastic Beanstalk for deploying the applications in your company. Unfortunately, you have discovered that there is a
resource type that is not supported by CloudFormation. What can you do to get around this.

A. Specify more mappings and separate your template into multiple templates by using nested stacks.
B. Create a custom resource type using template developer, custom resource template, and CloudFormatio
C. */
D. Specify the custom resource by separating your template into multiple templates by using nested stacks.
E. Use a configuration management tool such as Chef, Puppet, or Ansible.

Answer: B

Explanation:
Custom resources enable you to write custom provisioning logic in templates that AWS Cloud Formation runs anytime you create, update (if you changed the custom resource), or delete stacks. For example, you might want to include resources that aren't available as AWS Cloud Formation resource types. You can

include those resources by using custom resources. That way you can still manage all your related resources in a single stack.
For more information on custom resources in Cloudformation please visit the below URL:
? http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/template-custom- resources.htm I

NEW QUESTION 43
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 48
Your company has developed a web application and is hosting it in an Amazon S3 bucket configured for static website hosting. The application is using the AWS SDK for JavaScript in the browser to access data stored in an Amazon DynamoDB table. How can you ensure that API keys for access to your data in DynamoDB are kept secure?

A. Create an Amazon S3 role in 1AM with access to the specific DynamoDB tables, and assign it to the bucket hosting your website.
B. Configure S3 bucket tags with your AWS access keys for your bucket hosing your website so that the application can query them for access.
C. Configure a web identity federation role within 1AM to enable access to the correct DynamoDB resources and retrieve temporary credentials.
D. Store AWS keys in global variables within your application and configure the application to use these credentials when making requests.

Answer: C

Explanation:
With web identity federation, you don't need to create custom sign-in code or manage your own user identities. Instead, users of your app can sign in using a well-known identity provider (IdP) — such as Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible IdP, receive an authentication token, and then exchange that token for temporary security credentials in AWS that map to an 1AM role with permissions to use the resources in your AWS account. Using an IdP helps you keep your AWS account secure, because you don't have to embed and distribute long- term security credentials with your application. For more information on Web Identity Federation, please refer to the below document link: from AWS
http://docs.wsamazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html

NEW QUESTION 49
You need to create a Route53 record automatically in CloudFormation when not running in production during all launches of a Template. How should you implement this?

A. Use a Parameter for environment, and add a Condition on the Route53 Resource in the template to create the record only when environment is not production.
B. Create two templates, one with the Route53 record value and one with a null value for the recor
C. Use the one without it when deploying to production.
D. Use a Parameterfor environment, and add a Condition on the Route53 Resource in the template to create the record with a null string when environment is production.
E. Create two templates, one with the Route53 record and one without i

F. Use the one without it when deploying to production.

Answer: A

Explanation:
The optional Conditions section includes statements that define when a resource is created or when a property is defined. For example, you can compare whether a value is equal to another value. Based on the result of that condition, you can conditionally create resources. If you have multiple conditions, separate them with commas.
You might use conditions when you want to reuse a template that can create resources in different contexts, such as a test environment versus a production environment In your template, you can add an Environ me ntType input parameter, which accepts either prod or test as inputs. For the production environment, you might include Amazon CC2 instances with certain capabilities; however, for the test environment, you want to use reduced capabilities to save money. With conditions, you can define which resources are created and how they're configured for each environment type.
For more information on Cloudformation conditions please refer to the below link: http://docs.ws.amazon.com/AWSCIoudFormation/latest/UserGuide/cond itions-section- structure.htm I

NEW QUESTION 53
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 57
You are hired as the new head of operations for a SaaS company. Your CTO has asked you to make debugging any part of your entire operation simpler and as fast as possible. She complains that she has no idea what is going on in the complex, service-oriented architecture, because the developers just log to disk, and it's very hard to find errors in logs on so many services. How can you best meet this requirement and satisfy your CTO?

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

Answer: D

Explanation:
Amazon Dasticsearch Service makes it easy to deploy, operate, and scale dasticsearch for log analytics, full text search, application monitoring, and more. Amazon
Oasticsearch Service is a fully managed service that delivers Dasticsearch's easy-to-use APIs and real- time capabilities along with the availability, scalability, and security required by production workloads. The service offers built-in integrations with Kibana, Logstash, and AWS services including Amazon Kinesis Firehose, AWS Lambda, and Amazon Cloud Watch so that you can go from raw data to actionable insights quickly. For more information on Elastic Search, please refer to the below link:
• https://aws.amazon.com/elasticsearch-service/

NEW QUESTION 58
You have an application hosted in AWS. You wanted to ensure that when certain thresholds are reached, a Devops Engineer is notified. Choose 3 answers from the options given below

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

Answer: ACD

Explanation:
You can use Cloud Watch Logs to monitor applications and systems using log data. For example,
CloudWatch Logs can track the number of errors that occur in your
application logs and send you a notification whenever the rate of errors exceeds a threshold you specify. CloudWatch Logs uses your log data for monitoring; so, no code changes are required. For example, you can monitor application logs for specific literal terms (such as "NullReferenceLxception") or count the number of

occurrences of a literal term at a particular position in log data (such as "404" status codes in an Apache access log). When the term you are searching for is found, CloudWatch Logs reports the data to a CloudWatch metric that you specify. For more information on Cloudwatch Logs please refer to the below link:
http://docs.ws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html
Amazon CloudWatch uses Amazon SNS to send email. First, create and subscribe to an SNS topic.
When you create a CloudWatch alarm, you can add this SNS topic to send an email notification when the alarm changes state.
For more information on Cloudwatch and SNS please refer to the below link:
http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US_SetupSNS.html

NEW QUESTION 63
Your company releases new features with high frequency while demanding high application availability. As part of the application's A/B testing, logs from each updated Amazon EC2 instance of the application need to be analyzed in near real-time, to ensure that the application is working
flawlessly after each deployment. If the logs show any anomalous behavior, then the application version of the instance is changed to a more stable one. Which of the following methods should you use for shipping and analyzing the logs in a highly available manner?

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

Answer: C

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

NEW QUESTION 66
You need to deploy a Node.js application and do not have any experience with AWS. Which deployment method will be the simplest for you to deploy?

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

Answer: A

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

NEW QUESTION 68
Which of these is not an instrinsic function in AWS CloudFormation?

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

Answer: D

Explanation:
You can use intrinsic functions, such as Fn::If, Fn::Cquals, and Fn::Not, to conditionally create stack resources. These conditions are evaluated based on input parameters that you declare when you create or update a stack. After you define all your conditions, you can associate them with resources or resource properties in the Resources and Outputs sections of a template.
For more information on Cloud Formation template functions, please refer to the URL:
• http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function- reference.html and
• http://docs.aws.a mazon.com/AWSCloudFormation/latest/UserGuide/intri nsic-function- reference-conditions.html

NEW QUESTION 73
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 77
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 82
You are designing a service that aggregates clickstream data in batch and delivers reports to subscribers via email only once per week. Data is extremely spikey, geographically distributed, high- scale, and unpredictable. How should you design this system?

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

Answer: B

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

NEW QUESTION 83
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.

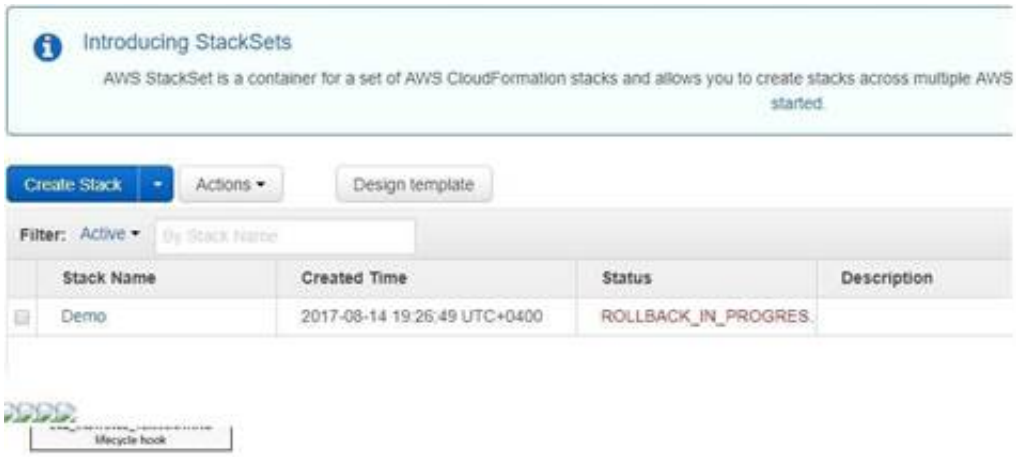

For more information on how Cloud Formation works, please refer to the below link: http://docs.ws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-whatis-howdoesitwork-html

NEW QUESTION 84
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 85
You need to scale an RDS deployment. You are operating at 10% writes and 90% reads, based on your logging. How best can you scale this in a simple way?

A. Create a second master RDS instance and peer the RDS groups.
B. Cache all the database responses on the read side with CloudFront.
C. Create read replicas for RDS since the load is mostly reads.
D. Create a Multi-AZ RDS installs and route read traffic to standby.

Answer: C

Explanation:
Amazon RDS Read Replicas provide enhanced performance and durability for database (DB) instances. This replication feature makes it easy to elastically scale out beyond the capacity constraints of a single DB Instance for read-heavy database workloads. You can create one or more replicas of a given source DB Instance and serve high-volume application read traffic from multiple copies of your data, thereby increasing aggregate read throughput. Read replicas can also be promoted when needed to become standalone DB instances.
Option A is invalid because you would need to maintain the synchronization yourself with a secondary instance.
Option B is invalid because you are introducing another layer unnecessarily when you already have read replica's Option D is invalid because you only use this for Standy's
For more information on Read Replica's, please refer to the below link: https://aws.amazon.com/rds/details/read-replicas/

NEW QUESTION 90
You are creating an application which stores extremely sensitive financial information. All information in the system must be encrypted at rest and in transit. Which of these is a violation of this policy?

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

Answer: A

Explanation:
If you use SSL termination, your servers will always get non-secure connections and will never know whether users used a more secure channel or not. If you are using Elastic beanstalk to configure the ELB, you can use the below article to ensure end to end encryption.
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/configuring-https-endtoend.html

NEW QUESTION 91
You are building a game high score table in DynamoDB. You will store each user's highest score for each game, with many games, all of which have relatively similar usage levels and numbers of players. You need to be able to look up the highest score for any game. What's the best DynamoDB key structure?

A. HighestScore as the hash/only key.

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

Answer: B

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

Partition key value	Uniformity
User ID, where the application has many users.	Good
Status code, where there are only a few possible status codes.	Bad
Item creation date, rounded to the nearest time period (e.g. day, hour, minute)	Bad
Device ID, where each device accesses data at relatively similar intervals	Good
Device ID, where even if there are a lot of devices being tracked, one is by far more popular than all the others.	Bad

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

NEW QUESTION 95
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 100
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 101
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 103
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 104
You need to create an audit log of all changes to customer banking data. You use DynamoDB to store this customer banking data. It's important not to lose any information due to server failures. What is an elegant way to accomplish this?

A. Use a DynamoDB StreamSpecification and stream all changes to AWS Lambd
B. Log the changes to AWS CloudWatch Logs, removing sensitive information before logging.
C. Before writing to DynamoDB, do a pre-write acknoledgment to disk on the application server, removing sensitive information before loggin
D. Periodically rotate these log files into S3.
E. Use a DynamoDB StreamSpecification and periodically flush to an EC2 instance store, removing sensitive information before putting the object
F. Periodically flush these batches to S3.
G. Before writing to DynamoDB, do a pre-write acknoledgment to disk on the application server, removing sensitive information before loggin
H. Periodically pipe these files into CloudWatch Logs.

Answer: A

Explanation:
You can use Lambda functions as triggers for your Amazon DynamoDB table. Triggers are custom actions you take in response to updates made to the DynamoDB table. To create a trigger, first you enable Amazon DynamoDB Streams for your table. Then, you write a Lambda function to process the updates published to the stream.
For more information on DynamoDB with Lambda, please visit the below URL: http://docs.aws.a mazon.com/lambda/latest/dg/with-ddb.html

NEW QUESTION 109
You need to deploy a new application version to production. Because the deployment is high-risk, you need to roll the new version out to users over a number of hours, to make sure everything is working correctly. You need to be able to control the proportion of users seeing the new version of the application down to the percentage point. You use ELB and EC2 with Auto Scaling Groups and custom AMIs with your code pre-installed assigned to Launch Configurations. There are no data base- level changes during your deployment. You have been told you cannot spend too much money, so you must not increase the number of EC2 instances much at all during the deployment, but you also need to be able to switch back to the original version of code quickly if something goes wrong. What is the best way to meet these requirements?

A. Create a second ELB, Auto Scaling Launch Configuration, and Auto Scaling Group using the Launch Configuratio
B. Create AMIs with all code pre-installe
C. Assign the new AMI to the second Auto Scaling Launch Configuratio
D. Use Route53 Weighted Round Robin Records to adjust the proportion of traffic hitting the two ELBs.S
E. Use the Blue-Green deployment method to enable the fastest possible rollback if neede
F. Create a full second stack of instances and cut the DNS over to the new stack of instances, and change the DNS back if a rollback is needed.
G. Create AMIs with all code pre-installe
H. Assign the new AMI to the Auto Scaling Launch Configuration, to replace the old on
I. Gradually terminate instances running the old code (launched with the old Launch Configuration) and allow the new AMIs to boot to adjust the traffic balance to the new cod
J. On rollback, reverse the process by doing the same thing, but changing the AMI on the Launch Config back to the original code.
K. Migrate to use AWS Elastic Beanstal
L. Use the established and well-tested Rolling Deployment setting AWS provides on the new Application Environment, publishing a zip bundle of the new code and

adjusting the wait period to spread the deployment over tim
M. Re-deploy the old code bundle to rollback if needed.

Answer: A

Explanation:
This is an example of a Blue Green Deployment
You can shift traffic all at once or you can do a weighted distribution. 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

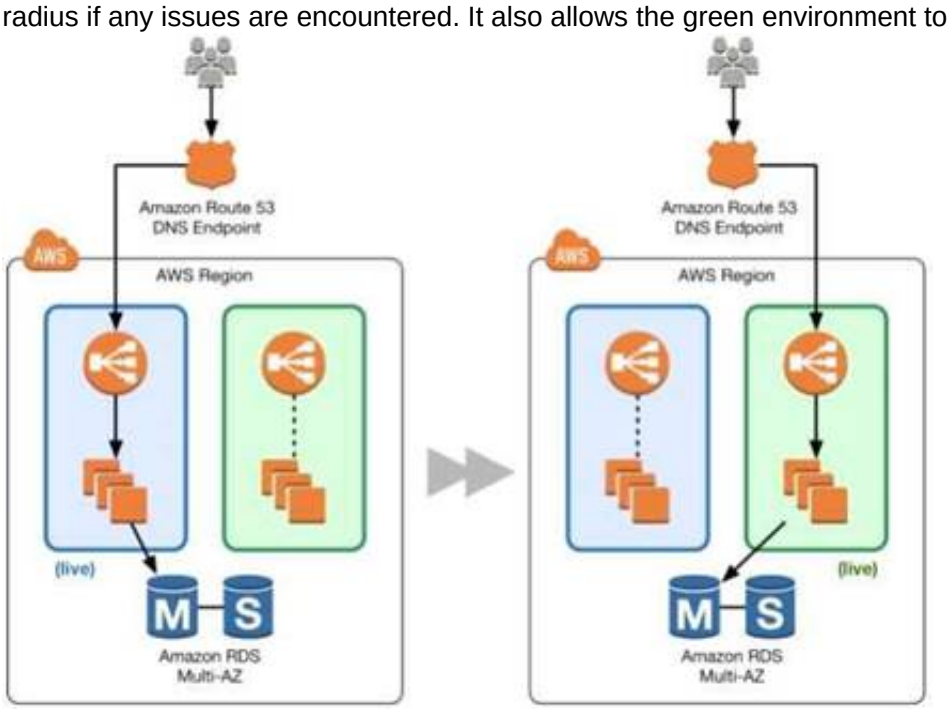

For more information on Blue Green Deployments, please visit the below URL:
• https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 112
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 115
Your team wants to begin practicing continuous delivery using CloudFormation, to enable automated builds and deploys of whole, versioned stacks or stack layers. You have a 3-tier, mission-critical system. Which of the following is NOT a best practice for using CloudFormation in a continuous delivery environment?

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

Answer: B

Explanation:
Answer - B
Some of the best practices for Cloudformation are
• Created Nested stacks
As your infrastructure grows, common patterns can emerge in which you declare the same components in each of your templates. You can separate out these common components and create dedicated templates for them. That way, you can mix and match different templates but use nested stacks to create a single, unified stack. Nested stacks are stacks that create other stacks. To create nested stacks, use the AWS::CloudFormation::Stackresource in your template to reference other templates.
• Reuse Templates
After you have your stacks and resources set up, you can reuse your templates to replicate your infrastructure in multiple environments. For example, you can create environments for development, testing, and production so that you can test changes before implementing them into production. To make templates reusable, use the parameters, mappings, and conditions sections so that you can customize your stacks when you create them. For example, for your

development environments, you can specify a lower-cost instance type compared to your production environment, but all other configurations and settings remain the same. For more information on Cloudformation best practises, please visit the below URL:
http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/best-practices.html

NEW QUESTION 118
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 121
You need to deploy an AWS stack in a repeatable manner across multiple environments. You have selected CloudFormation as the right tool to accomplish this, but have found that there is a resource type you need to create and model, but is unsupported by CloudFormation. How should you overcome this challenge?

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

Answer: D

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

NEW QUESTION 123
Your API requires the ability to stay online during AWS regional failures. Your API does not store any state, it only aggregates data from other sources - you do not have a database. What is a simple but effective way to achieve this uptime goal?

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

Answer: D

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

NEW QUESTION 125
Your serverless architecture using AWS API Gateway, AWS Lambda, and AWS DynamoDB experienced a large increase in traffic to a sustained 3000 requests per second, and dramatically increased in failure rates. Your requests, during normal operation, last 500 milliseconds on average. Your DynamoDB table did not

exceed 50% of provisioned throughput, and Table primary keys are designed correctly. What is the most likely issue?

A. Your API Gateway deployment is throttling your requests.
B. Your AWS API Gateway Deployment is bottleneckingon request (deserialization.
C. You did not request a limit increase on concurrent Lambda function executions.
D. You used Consistent Read requests on DynamoDB and are experiencing semaphore lock.

Answer: C

Explanation:
Every Lambda function is allocated with a fixed amount of specific resources regardless of the memory allocation, and each function is allocated with a fixed amount of code storage per function and per account.
By default, AWS Lambda limits the total concurrent executions across all functions within a given region to 1000.
For more information on Concurrent executions, please visit the below URL: http://docs.aws.amazon.com/lambda/latest/dg/concurrent-executions.html

NEW QUESTION 126
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
You need your CI to build AM Is with code pre-installed on the images on every new code push. You need to do this as cheaply as possible. How do you do this?

A. Bid on spot instances just above the asking price as soon as new commits come in, perform all instance configuration and setup, then create an AM I based on the spot instance.
B. Have the CI launch a new on-demand EC2 instance when new commits come in, perform all instance configuration and setup, then create an AMI based on the on-demand instance.
C. Purchase a Light Utilization Reserved Instance to save money on the continuous integration machin
D. Use these credits whenever your create AMIs on instances.
E. When the CI instance receives commits, attach a new EBS volume to the CI machin
F. Perform all setup on this EBS volume so you don't need

Answer: A

Explanation:
Amazon CC2 Spot instances allow you to bid on spare Amazon CC2 computing capacity. Since Spot instances are often available at a discount compared to On-Demand pricing, you can significantly reduce the cost of running your applications, grow your application's compute capacity and throughput for the same budget, and enable new types of cloud computing applications.
For more information on Spot Instances, please visit the below URL: https://aws.amazon.com/ec2/spot/

NEW QUESTION 130
An EC2 instance has failed a health check. What will the ELB do?

A. The ELB will terminate the instance
B. The ELB stops sending traffic to the instance that failed its health check
C. The ELB does nothing
D. The ELB will replace the instance

Answer: B

Explanation:
The AWS Documentation mentions
The load balancer routes requests only to the healthy instances. When the load balancer determines that an instance is unhealthy, it stops routing requests to that instance. The load balancer resumes routing requests to the instance when it has been restored to a healthy state.
For more information on ELB health checks, please refer to the below link: http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-healthchecks.html

NEW QUESTION 132
Which of the following services can be used in conjunction with Cloudwatch Logs. Choose the 3 most viable services from the options given below

A. Amazon Kinesis
B. Amazon S3
C. Amazon SQS

D. Amazon Lambda

Answer: ABD

Explanation:
The AWS Documentation the following products which can be integrated with Cloudwatch logs
1) Amazon Kinesis - Here data can be fed for real time analysis
2) Amazon S3 - You can use CloudWatch Logs to store your log data in highly durable storage such as S3.
3) Amazon Lambda - Lambda functions can be designed to work with Cloudwatch log For more information on Cloudwatch Logs, please refer to the below link:
link:http://docs^ws.amazon.com/AmazonCloudWatch/latest/logs/WhatlsCloudWatchLogs.html

NEW QUESTION 133
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 135
There is a requirement for a vendor to have access to an S3 bucket in your account. The vendor already has an AWS account. How can you provide access to the vendor on this bucket.

A. Create a new 1AM user and grant the relevant access to the vendor on that bucket.
B. Create a new 1AM group and grant the relevant access to the vendor on that bucket.
C. Create a cross-account role for the vendor account and grant that role access to the S3 bucket.
D. Create an S3 bucket policy that allows the vendor to read from the bucket from their AWS account.

Answer: C

Explanation:
The AWS documentation mentions
You share resources in one account with users in a different account. By setting up cross-account access in this way, you don't need to create individual 1AM users in each account In addition, users don't have to sign out of one account and sign into another in order to access resources that are in different AWS accounts. After configuring the role, you see how to use the role from the AWS Management Console, the AWS CLI, and the API
For more information on Cross Account Roles Access, please refer to the below link:
• http://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 140
Which of the following is the right sequence of initial steps in the deployment of application revisions using Code Deploy
1) Specify deployment configuration
2) Upload revision
3) Create application
4) Specify deployment group

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

Answer: C

Explanation:
The below diagram from the AWS documentation shows the deployment steps

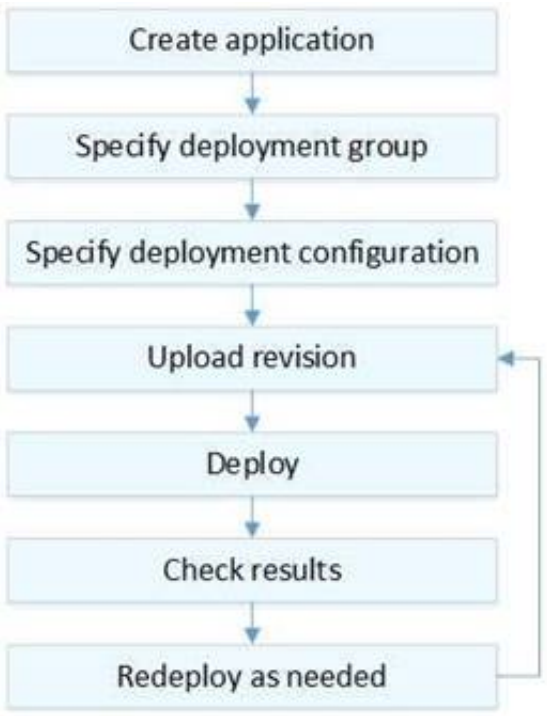

For more information on the deployment steps please refer to the below link:
• http://docs.aws.amazon.com/codedeploy/latest/userguide/de ployment-steps.html

NEW QUESTION 144
You need to investigate one of the instances which is part of your Autoscaling Group. How would you implement this.

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

Answer: B

Explanation:
The AWS Documentation mentions
Auto Scaling enables you to put an instance that is in the InService state into the Standbystate, update or troubleshoot the instance, and then return the instance to service. Instances that are on standby are still part of the Auto Scaling group, but they do not actively handle application traffic.
For more information on the standby state please refer to the below link:
• http://docs.aws.a mazon.com/autoscaling/latest/userguide/as-enter-exit-standby.html

NEW QUESTION 149
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 150
You need to deploy a multi-container Docker environment on to Elastic beanstalk. Which of the following files can be used to deploy a set of Docker containers to Elastic beanstalk

A. Dockerfile
B. DockerMultifile
C. Dockerrun.aws.json
D. Dockerrun

Answer: C

Explanation:
The AWS Documentation specifies
A Dockerrun.aws.json file is an 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.
For more information on this, please visit the below URL:
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_v2config.html

NEW QUESTION 152
Which of the following is not a rolling type update which is present for Configuration Updates when it comes to the Elastic Beanstalk service

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

Answer: B

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

The AWS Documentation mentions
1) With health-based rolling updates. Elastic Beanstalk waits until instances in a batch pass health checks before moving on to the next batch.
2) For time-based rolling updates, you can configure the amount of time that Elastic Beanstalk waits after completing the launch of a batch of instances before moving on to the next batch. This pause time allows your application to bootsrap and start serving requests.
3) Immutable environment updates are an alternative to rolling updates that ensure that configuration changes that require replacing instances are applied efficiently and safely. If an immutable environment update fails, the rollback process requires only terminating an Auto Scalinggroup. A failed rolling update, on the other hand, requires performing an additional rolling update to roll back the changes.
For more information on Rolling updates for Elastic beanstalk configuration updates, please visit the below URL:
• http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.ro11ingupdates.html

NEW QUESTION 153
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 154
Your company has a set of resources hosted in AWS. They want to be notified when the costs of the AWS resources running in the account reaches a certain threshold. How can this be accomplished in an ideal way.

A. Create a script which monitors all the running resources and calculates the costs accordingly.
B. Download the cost reports and analyze the reports to see if the costs are going beyond the threshold
C. Create a billing alarm which can alert you when the costs are going beyond a certain threshold
D. Create a consolidated billing report and see if the costs are going beyond the threshold.

Answer: C

Explanation:
The AWS Documentation mentions
You can monitor your AWS costs by using Cloud Watch. With Cloud Watch, you can create billing alerts that notify you when your usage of your services exceeds thresholds that you define. You specify these threshold amounts when you create the billing alerts.
When your usage exceeds these amounts, AWS sends you an email notification. You can also sign up to receive notifications when AWS prices change. For more

information on billing alarms, please visit the below URL:
• http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/monitor-charges.html

NEW QUESTION 157
Your company has a set of resources hosted in AWS. Your IT Supervisor is concerned with the costs being incurred by the resources running in AWS and wants to optimize on the costs as much as possible. Which of the following ways could help achieve this efficiently? Choose 2 answers from the options given below.

A. Create Cloudwatch alarms to monitor underutilized resources and either shutdown or terminate resources which are not required.
B. Use the Trusted Advisor to see underutilized resources
C. Create a script which monitors all the running resources and calculates the costs accordingl
D. The analyze those resources accordingly and see which can be optimized.
E. Create Cloudwatch logs to monitor underutilized resources and either shutdown or terminate resources which are not required.

Answer: AB

Explanation:
You can use Cloudwatch alarms to see if resources are below a threshold for long periods of time. If so you can take the decision to either stop them or to terminate the resources.
For more information on Cloudwatch alarms, please visit the below URL:
• <http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/Ala rmThatSendsCmail.html>
In the Trusted Advisor, when you enable the Cost optimization section, you will get all sorts of checks which can be used to optimize the costs of your AWS resources.

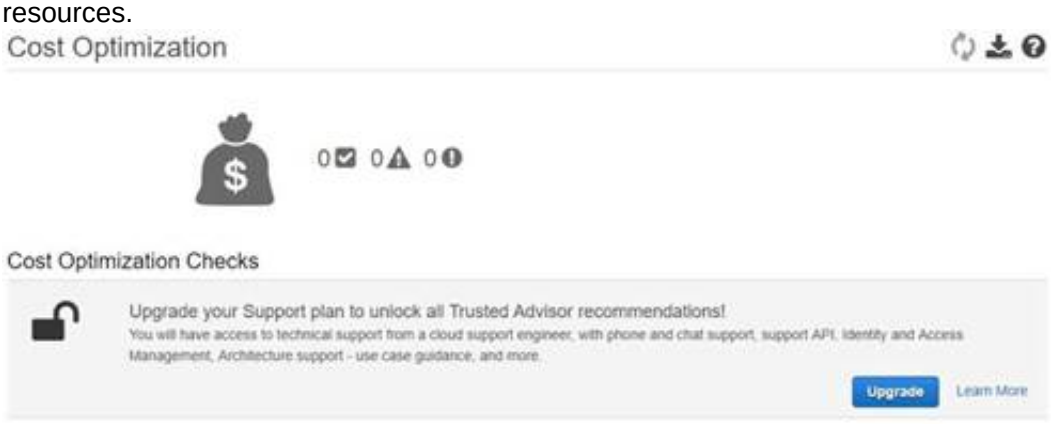

For more information on the Trusted Advisor, please visit the below U RL:
• https://aws.amazon.com/premiumsupport/trustedadvisor/

NEW QUESTION 161
An audit is going to be conducted for your company's AWS account. Which of the following steps will ensure that the auditor has the right access to the logs of your AWS account

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

Answer: B

Explanation:

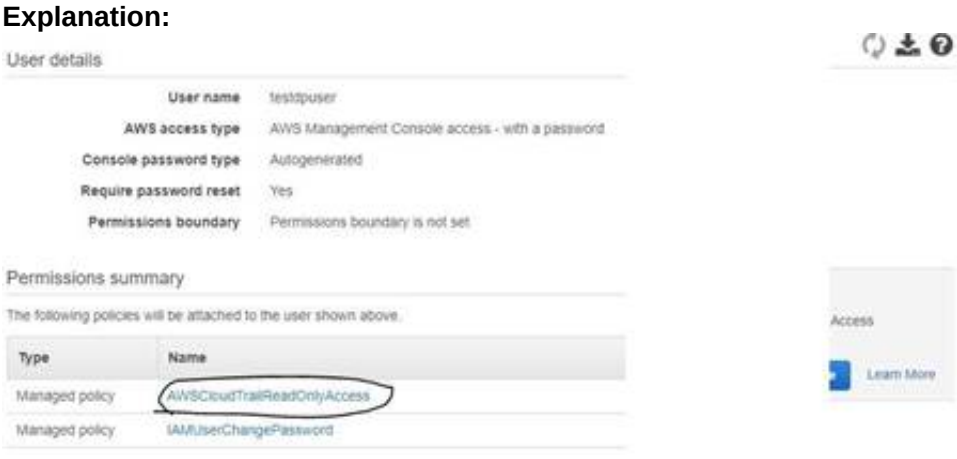

The AWS Documentation clearly mentions the below
AWS CloudTrail is an AWS service that helps you enable governance, compliance, and operational and risk auditing of your AWS account. Actions taken by a user,
role, or an AWS service are recorded as events in CloudTrail. Events include actions taken in the AWS Management Console, AWS Command Line Interface, and AWS SDKs and APIs.
For more information on Cloudtrail, please visit the below URL:
• http://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html

NEW QUESTION 162
Your company has a set of EC2 resources hosted on AWS. Your new IT procedures state that AWS EC2 Instances must be of a particular Instance type. Which of the following can be used to get the list of EC2 Instances which currently don't match the instance type specified in the new IT procedures

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

Answer: B

Explanation:
In AWS Config, you can create a rule which can be used to check if CC2 Instances follow a particular instance type. Below is a snapshot of the output of a rule to check if CC2 instances matches the type of t2micro.

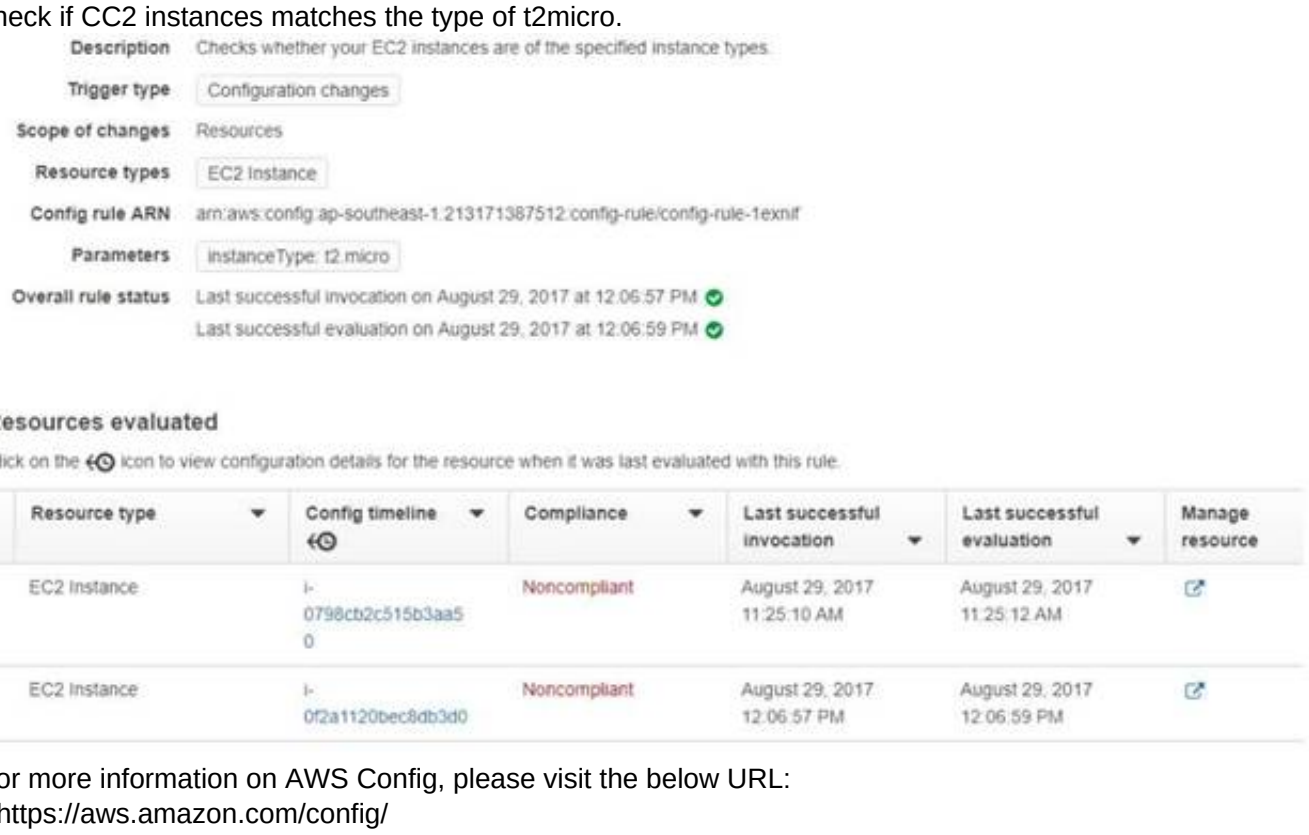

For more information on AWS Config, please visit the below URL:
• https://aws.amazon.com/config/

NEW QUESTION 164
Your company is planning to setup a wordpress application. The wordpress application will connect to a MySQL database. Part of the requirement is to ensure that the database environment is fault
tolerant and highly available. Which of the following 2 options individually can help fulfil this requirement.

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

Answer: AD

Explanation:
One way to ensure high availability and fault tolerant environments is to ensure Instances are located across multiple availability zones. Hence if you are hosting MySQL yourself, ensure you have instances spread across multiple AZ's
The AWS Documentation mentions the following about the multi-AZ feature
Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. Amazon RDS uses several different technologies to provide failover support. Multi-AZ deployments for Oracle, PostgreSGL, MySQL, and MariaDB DB instances use Amazon's failover technology
For more information on AWS Multi-AZ deployments, please visit the below URL:
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html

NEW QUESTION 165
You are currently using Elastic Beanstalk to host your production environment. You need to rollout updates to your application hosted on this environment. This is a critical application which is why there is a requirement that the rollback, if required, should be carried out with the least amount of downtime. Which of the following deployment strategies would ideally help achieve this purpose

A. Create a Cloudformation template with the same resources as those in the Elastic beanstalk environmen
B. If the deployment fails, deploy the Cloudformation template.
C. Use Rolling updates in Elastic Beanstalk so that if the deployment fails, the rolling updates feature would roll back to the last deployment.
D. Create another parallel environment in elastic beanstal
E. Use the Swap URL feature.
F. Create another parallel environment in elastic beanstal
G. Create a new Route53 Domain name for the new environment and release that url to the users.

Answer: C

Explanation:
Since the requirement is to have the least amount of downtime, the ideal way is to create a blue green deployment environment and then use the Swap URL feature
to swap environments for the new deployment and then do the swap back, incase the deployment fails.
The AWS Documentation mentions the following on the SWAP url feature of Elastic Beanstalk
Because Elastic Beanstalk performs an in-place update when you update your application versions, your application may become unavailable to users for a short period of time. It is possible to avoid this downtime by performing a blue/green deployment, where you deploy the new version to a separate environment, and then swap CNAMCs of the two environments to redirect traffic to the new version instantly.

NEW QUESTION 170
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.html

NEW QUESTION 174
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 175
The company you work for has a huge amount of infrastructure built on AWS. However there has been some concerns recently about the security of this infrastructure, and an external auditor has been given the task of running a thorough check of all of your company's AWS assets. The auditor will be in the USA while your company's infrastructure resides in the Asia Pacific (Sydney) region on AWS. Initially, he needs to check all of your VPC assets, specifically, security groups and NACLs You have been assigned the task of providing the auditor with a login to be able to do this. Which of the following would be the best and most secure solution to provide the auditor with so he can begin his initial investigations? Choose the correct answer from the options below

A. Create an 1AM user tied to an administrator rol
B. Also provide an additional level of security with MFA.
C. Give him root access to your AWS Infrastructure, because he is an auditor he will need access to every service.
D. Create an 1AM user who will have read-only access to your AWS VPC infrastructure and provide the auditor with those credentials.
E. Create an 1AM user with full VPC access but set a condition that will not allow him to modify anything if the request is from any IP other than his own.

Answer: C

Explanation:
Generally you should refrain from giving high level permissions and give only the required permissions. In this case option C fits well by just providing the relevant access which is required.
For more information on 1AM please see the below link:
• https://aws.amazon.com/iam/

NEW QUESTION 179
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 180
You are incharge of creating a Cloudformation template that will be used to spin our resources on demand for your Devops team. The requirement is that this cloudformation template should be able to spin up resources in different regions. Which of the following aspects of Cloudformation templates can help you design the template to spin up resources based on the region.

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

Answer: A

Explanation:
The AWS Documentation mentions
The optional Mappings section matches a key to a corresponding set of named values. For example, if you want to set values based on a region, you can create a mapping that uses the region name as a key and contains the values you want to specify for each specific region. You use the Fn::FindInMap intrinsic function to retrieve values in a map.
For more information on mappings please refer to the below link:
? http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/mappings-section-structure.html

NEW QUESTION 184
You are the IT administrator for your company. You have the responsibility of creating development environments which would confirm to the LAMP development stack. The requirement is that the
development team always gets the latest version of the LAMP stack each time a new instance is launched. Which of the following is an efficient and effective way to implement this requirement? Choose 2 answers from the options given below

A. Create an AMI with all the artifacts of the LAMP stack and provide an instance to the development team based on the AMI.
B. Create a cloudformation template and use the cloud-init directives to download and the install the LAMP stack packages.
C. Use the User data section and use a custom script which will be used to download the necessary LAMP stack packages.
D. Create an EBS Volume with the LAMP stack and attach it to an instance whenever it is required.

Answer: BC

Explanation:
Using User data and cloud-init directives you can always ensure you download the latest version of the LAMP stack and give it to the development teams. With AMI's
you will always have the same version and will need to create an AMI everytime the version of the LAMP stack changes.
The AWS Documentation mentions
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 User data please refer to the below link: http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/user-data.html

NEW QUESTION 187
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 188
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 193
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.

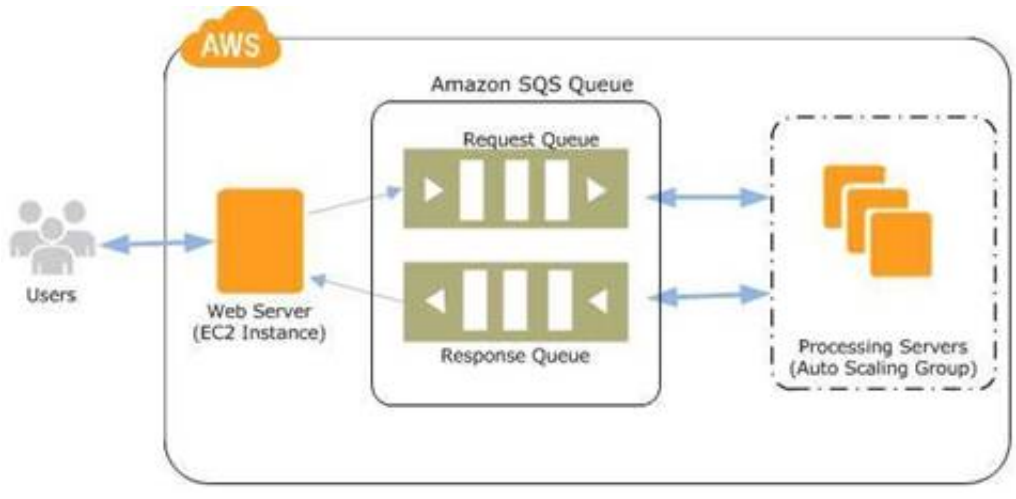

For more information on SQS and Autoscaling, please refer to the below U RL: http://docs.aws.amazon.com/autoscaling/latest/userguide/as-using-sqs-queue.html

NEW QUESTION 198
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

DOP-C01 Practice Exam Features:

* DOP-C01 Questions and Answers Updated Frequently

* DOP-C01 Practice Questions Verified by Expert Senior Certified Staff

* DOP-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* DOP-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year