Exam Questions DOP-C01

AWS Certified DevOps Engineer- Professional

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

NEW QUESTION 1
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 2
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

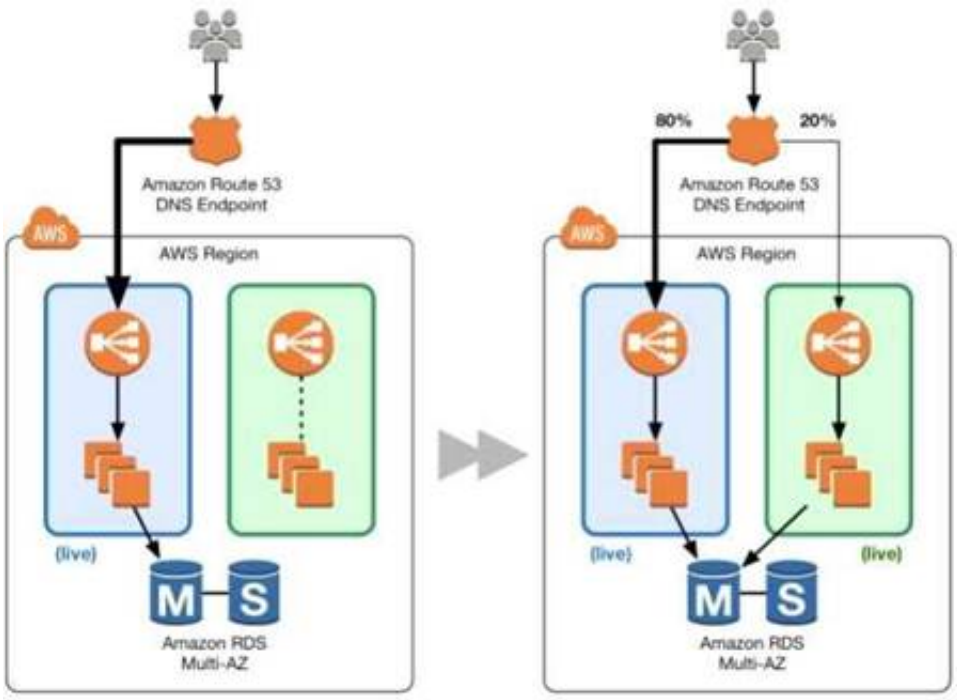

1) First create a new ELB which will be used to point to the new production changes.
2) Use the Weighted Route policy for Route53 to distribute the traffic to the 2 ELB's based on a 80- 20% traffic scenario. This is the normal case, the % can be changed based on the requirement.
3) Finally when all changes 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 3
You have just recently deployed an application on EC2 instances behind an ELB. After a couple of weeks, customers are complaining on receiving errors from the application. You want to diagnose the errors and are trying to get errors from the ELB access logs. But the ELB access logs are empty. What is the reason for this.

A. You do not have the appropriate permissions to access the logs
B. You do not have your CloudWatch metrics correctly configured
C. ELB Access logs are only available for a maximum of one week.
D. Access logging is an optional feature of Elastic Load Balancing that is disabled by default

Answer: D

Explanation:
Clastic Load Balancing provides access logs that capture detailed information about requests sent to
your load balancer. Cach log contains information such as the
time the request was received, the client's IP address, latencies, request paths, and server responses.
You can use these access logs to analyze traffic patterns and to troubleshoot issues.
Access logging is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logging for your load balancer. Clastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify. You can disable access logging at any time.
For more information on CLB access logs, please refer to the below document link: from AWS
http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/access-log-collection.html

NEW QUESTION 4
You have an ELB setup in AWS with EC2 instances running behind it. You have been requested to monitor the incoming connections to the ELB. Which of the below options can suffice this requirement?

A. UseAWSCIoudTrail with your load balancer
B. Enable access logs on the load balancer
C. Use a CloudWatch Logs Agent
D. Create a custom metric CloudWatch filter on your load balancer

Answer: B

Explanation:
Clastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Cach log contains information such as the
time the request was received, the client's IP address, latencies, request paths, and server responses.
You can use these access logs to analyze traffic patterns and to troubleshoot issues.
Option A is invalid because this service will monitor all AWS services Option C and D are invalid since CLB already provides a logging feature.
For more information on ELB access logs, please refer to the below document link: from AWS
http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/access-log-collection.html

NEW QUESTION 5
You are responsible for your company's large multi-tiered Windows-based web application running on Amazon EC2 instances situated behind a load balancer. While reviewing metrics, you've started noticing an upwards trend for slow customer page load time. Your manager has asked you to come up with a solution to ensure that customer load time is not affected by too many requests per second. Which technique would you use to solve this issue?

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

Answer: C

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

NEW QUESTION 6
During metric analysis, your team has determined that the company's website during peak hours is experiencing response times higher than anticipated. You currently rely on Auto Scaling to make sure that you are scaling your environment during peak windows. How can you improve your Auto Scaling policy to reduce this high response time? Choose 2 answers.

A. Push custom metrics to CloudWatch to monitor your CPU and network bandwidth from your servers, which will allow your Auto Scaling policy to have betterfine-grain insight.
B. IncreaseyourAutoScalinggroup'snumberofmaxservers.
C. Create a script that runs and monitors your servers; when it detects an anomaly in load, it posts to an Amazon SNS topic that triggers Elastic Load Balancing to add more servers to the load balancer.
D. Push custom metrics to CloudWatch for your application that include more detailed information about your web application, such as how many requests it is handling and how many are waiting to be processed.

Answer: BD

Explanation:
Option B makes sense because maybe the max servers is low hence the application cannot handle the peak load.
Option D helps in ensuring Autoscaling can scale the group on the right metrics.
For more information on Autoscaling health checks, please refer to the below document link: from AWS
http://docs.aws.amazon.com/autoscaling/latest/userguide/healthcheck.html

NEW QUESTION 7
You currently run your infrastructure on Amazon EC2 instances behind an Auto Scalinggroup. All logs for your application are currentl\ written to ephemeral storage. Recently your company experienced a major bug in the code that made it through testing and was ultimately deployed to your fleet. This bug triggered your Auto Scalinggroup to scale up and back down before you could successfully retrieve the logs off your server to better assist you in troubleshooting the bug. Which technique should you use to make sure you are able to review your logs after your instances have shut down?

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

Answer: C

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

NEW QUESTION 8
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 9
You have been requested to use CloudFormation to maintain version control and achieve automation for the applications in your organization. How can you best use CloudFormation to keep everything agile and maintain multiple environments while keeping cost down?

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

Answer: A

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

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

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

Answer: B

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

NEW QUESTION 15
You have a complex system that involves networking, 1AM policies, and multiple, three-tier applications. You are still receiving requirements for the new system, so you don't yet know how many AWS components will be present in the final design. You want to start using AWS CloudFormation to define these AWS resources so that you can automate and version-control your infrastructure. How would you use AWS CloudFormation to provide agile new environments for your customers in a cost-effective, reliable manner?

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

Answer: B

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

NEW QUESTION 20
You use Amazon Cloud Watch as your primary monitoring system for your web application. After a
recent software deployment, your users are getting Intermittent 500 Internal Server Errors when using the web application. You want to create a Cloud Watch alarm, and notify an on-call engineer when these occur. How can you accomplish this using AWS services? Choose three answers from the options given below

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

Answer: BDE

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

NEW QUESTION 21
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
Your development team wants account-level access to production instances in order to do live debugging of a highly secure environment. Which of the following should you do?

A. Place the credentials provided by Amazon Elastic Compute Cloud (EC2) into a secure Amazon Sample Storage Service (S3) bucket with encryption enable
B. Assign AWS Identity and Access Management (1AM) users to each developer so they can download the credentials file.
C. Place an internally created private key into a secure S3 bucket with server-side encryption using customer keys andconfiguration management, create a service account on al I the instances using this private key, and assign I AM users to each developer so they can download the fi le.
D. Place each developer's own public key into a private S3 bucket, use instance profiles and configuration management to create a user account for each developer on all instances, and place the user's public keys into the appropriate accoun
E. ^/
F. Place the credentials provided by Amazon EC2 onto an MFA encrypted USB drive, and physically share it with each developer so that the private key never leaves the office.

Answer: C

Explanation:
An instance profile is a container for an 1AM role that you can use to pass role information to an CC2 instance when the instance starts.
A private S3 bucket can be created for each developer, the keys can be stored in the bucket and then assigned to the instance profile.
Option A and D are invalid, because the credentials should not be provided by a AWS EC2 Instance. Option B is invalid because you would not create a service account, instead you should create an instance profile.
For more information on Instance profiles, please refer to the below document link: from AWS
• http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-ro le-ec2_instance- profiles.htm I

NEW QUESTION 26
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

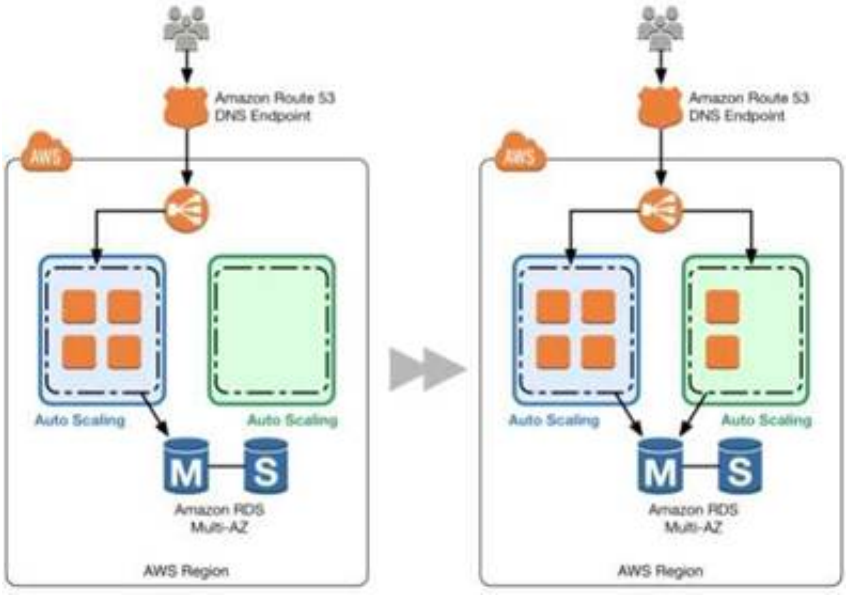

A blue group carries the production load while a green group is staged and deployed with the new code. When if s time to deploy, you simply attach the green group to
the existing load balancer to introduce traffic to the new environment. For HTTP/HTTP'S listeners, the load balancer favors the green Auto 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 31
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 36
Your application uses Cloud Formation to orchestrate your application's resources. During your testing phase before the application went live, your Amazon RDS instance type was changed and caused the instance to be re-created, resulting In the loss of test data. How should you prevent this from occurring in the future?

A. Within the AWS CloudFormation parameter with which users can select the Amazon RDS instance type, set AllowedValues to only contain the current instance

type.
B. Use an AWS CloudFormation stack policy to deny updates to the instanc
C. Only allow UpdateStack permission to 1AM principals that are denied SetStackPolicy.
D. In the AWS CloudFormation template, set the AWS::RDS::DBInstance's DBInstanceClass property to be read-only.
E. Subscribe to the AWS CloudFormation notification "BeforeResourcellpdate," and call CancelStackUpdate if the resource identified is the Amazon RDS instance.
F. Update the stack using ChangeSets

Answer: E

Explanation:
When you need to update a stack, understanding how your changes will affect running resources before you implement them can help you update stacks with confidence. Change sets allow you to preview how proposed changes to a stack might impact your running resources, for example, whether your changes will delete or replace any critical resources, AWS CloudFormation makes the changes to your stack only when you decide to execute the change set, allowing you to decide whether to proceed with your proposed changes or explore other changes by creating another change set
For example, you can use a change set to verify that AWS CloudFormation won't replace your stack's database instances during an update.

NEW QUESTION 38
You have a set of EC2 instances hosted in AWS. You have created a role named DemoRole and assigned that role to a policy, but you are unable to use that role with an instance. Why is this the case.

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

Answer: A

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

NEW QUESTION 40
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 need to monitor specific metrics from your application and send real-time alerts to your Devops Engineer. Which of the below services will fulfil this requirement? Choose two answers

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

Answer: AB

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

NEW QUESTION 48
You have an application hosted in AWS. This application was created using Cloudformation Templates and Autoscaling. Now your application has got a surge of users which is decreasing the performance of the application. As per your analysis, a change in the instance type to C3 would resolve the issue. Which of the below option can introduce this change while minimizing downtime for end users?

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

Answer: B

Explanation:
Ensure first that the cloudformation template is updated with the new instance type.
The AWS::AutoScaling::AutoScalingGroup resource supports an UpdatePolicy attribute. This is used to define how an Auto Scalinggroup resource is updated when
an update to the Cloud Formation stack occurs. A common approach to updating an Auto Scaling group is to perform a rolling update, which is done by specifying the AutoScalingRollingUpdate policy. This retains the same Auto Scaling group and replaces old instances with new ones, according to the parameters specified.
Option A is invalid because this will cause an interruption to the users.
Option C is partially correct, but it does not have all the steps as mentioned in option B.
Option D is partially correct, but we need the AutoScalingRollingUpdate attribute to ensure a rolling update is peformed.
For more information on AutoScaling Rolling updates please refer to the below link:
• https://aws.amazon.com/premiumsupport/knowledge-center/auto-scaling-group-rolling- updates/

NEW QUESTION 51
What is web identity federation?

A. Use of an identity provider like Google or Facebook to become an AWS1AM User.
B. Use of an identity provider like Google or Facebook to exchange for temporary AWS security credentials.
C. Use of AWS 1AM Usertokens to log in as a Google or Facebook user.
D. Use STS service to create an user on AWS which will allow them to login from facebook orgoogle app.

Answer: B

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 link:
http://docs^ws.amazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html

NEW QUESTION 52
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 53
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 55
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 59
You are planning on using the Amazon RDS facility for Fault tolerance for your application. How does Amazon RDSmuIti Availability Zone model work

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

Answer: A

Explanation:
Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB)
Instances, making them a natural fit for production database
workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Cach AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable.
In case of an infrastructure failure, Amazon RDS performs an automatic failover to the standby (or to a read replica in the case of Amazon Aurora), so that you can resume database operations as soon as the failover is complete.
The below diagram from the AWS documentation shows how this is configured

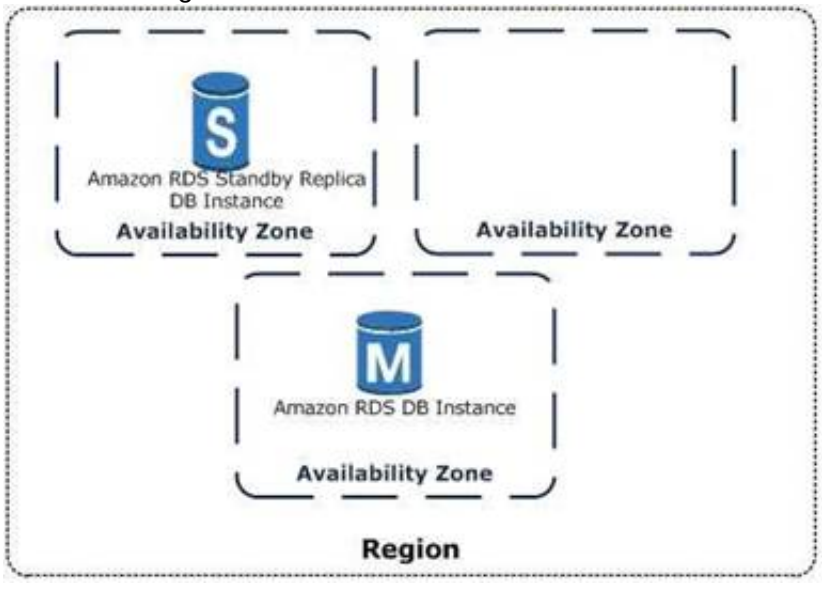

Option B is invalid because the replication is synchronous.
Option C and D are invalid because this is built around AZ and not regions. For more information on Multi-AZ RDS, please visit the below URL:
https://aws.amazon.com/rds/details/multi-az/

NEW QUESTION 62
You have a requirement to host a cluster of NoSQL databases. There is an expectation that there will be a lot of I/O on these databases. Which EBS volume type is best for high performance NoSQL cluster deployments?

A. io1

B. gp1
C. standard
D. gp2

Answer: A

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

NEW QUESTION 63
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 67
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 68
You are planning on using encrypted snapshots in the design of your AWS Infrastructure. Which of the following statements are true with regards to EBS Encryption

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

Answer: C

Explanation:
Amazon CBS encryption offers you a simple encryption solution for your CBS volumes without the need for you to build, maintain, and secure your own key management infrastructure. When you create an encrypted CBS volume and attach it to a supported instance type, the following types of data are encrypted:
• Data at rest inside the volume
• All data moving between the volume and the instance
• All snapshots created from the volume
Snapshots that are taken from encrypted volumes are automatically encrypted. Volumes that are created from encrypted snapshots are also automatically encrypted.
For more information on CBS encryption, please visit the below URL:
• http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/ CBSCncryption.html

NEW QUESTION 69
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 71
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 76
What is required to achieve gigabit network throughput on EC2? You already selected cluster- compute, 10GB instances with enhanced networking, and your workload is already network-bound, but you are not seeing 10 gigabit speeds.

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

Answer: D

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

NEW QUESTION 79
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 81
Your CTO thinks your AWS account was hacked. What is the only way to know for certain if there was unauthorized access and what they did, assuming your hackers are very sophisticated AWS engineers and doing everything they can to cover their tracks?

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

Answer: A

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

NEW QUESTION 86
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 87
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 89
When creating an Elastic Beanstalk environment using the Wizard, what are the 3 configuration options presented to you

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

Answer: ABD

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

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

NEW QUESTION 94
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 99
You have carried out a deployment using Elastic Beanstalk with All at once method, but the application is unavailable. What could be the reason for this

A. You need to configure ELB along with Elastic Beanstalk
B. You need to configure Route53 along with Elastic Beanstalk
C. There will always be a few seconds of downtime before the application is available
D. The cooldown period is not properly configured for Elastic Beanstalk

Answer: C

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

NEW QUESTION 104
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 106
Which of the following credentials types are supported by AWSCodeCommit? Select 3 Options

A. Git Credentials
B. SSH Keys
C. User name/password
D. AWS Access Kevs

Answer: ABD

Explanation:
The AWS documentation mentions
I AM supports AWS CodeCommit with three types of credentials:
Git credentials, an 1AM -generated user name and password pair you can use to communicate with AWS CodeCommit repositories over HTTPS.
SSH keys, a locally generated public-private key pair that you can associate with your 1AM user to communicate with AWS CodeCommit repositories over SSH.
AWS access keys, which you can use with the credential helper included with the AWS CLI to communicate with AWS CodeCommit repositories over HTTPS.
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_ssh-keys.html

NEW QUESTION 111
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

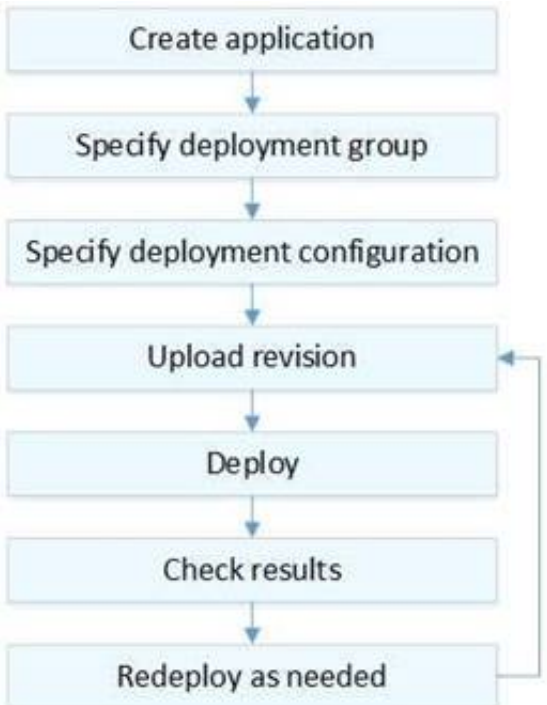

For more information on the deployment steps please refer to the below link:
• http://docs.aws.amazon.com/codedeploy/latest/userguide/de ployment-steps.html

NEW QUESTION 116
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 120
You have a requirement to automate the creation of EBS Snapshots. Which of the following can be used to achieve this in the best way possible.

A. Create a powershell script which uses the AWS CLI to get the volumes and then run the script as a cron job.

B. Use the AWSConfig service to create a snapshot of the AWS Volumes
C. Use the AWS CodeDeploy service to create a snapshot of the AWS Volumes
D. Use Cloudwatch Events to trigger the snapshots of EBS Volumes

Answer: D

Explanation:
The best is to use the inbuilt service from Cloudwatch, as Cloud watch Events to automate the creation of CBS Snapshots. With Option A, you would be restricted to
running the powrshell script on Windows machines and maintaining the script itself And then you have the overhead of having a separate instance just to run that script.
When you go to Cloudwatch events, you can use the Target as EC2 CreateSnapshot API call as shown below.
Create rules to invoke Targets based on Events happening in your AWS environment.

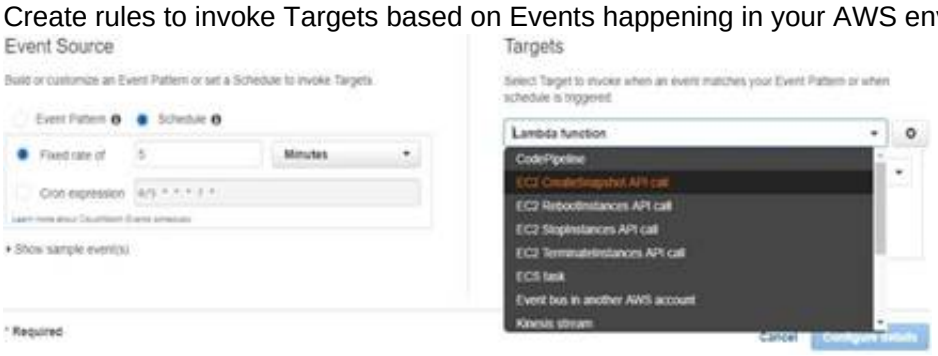

The AWS Documentation mentions
Amazon Cloud Watch Cvents delivers a near real-time stream of system events that describe changes in Amazon Web Services (AWS) resources. Using simple rules
that you can quickly set up, you can match events and route them to one or more target functions or streams. Cloud 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 Cvents, please visit the below U RL:
• http://docs.aws.amazon.com/AmazonCloudWatch/latest/events/Whatl sCloudWatchCvents.html

NEW QUESTION 121
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 123
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 124
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 127
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::FindlnMap intrinsic function to retrieve values in a map.
For more information on mappings please refer to the below link:
? http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/mappings-section-structure.html

NEW QUESTION 128
You are working as an AWS Devops admins for your company. You are in-charge of building the infrastructure for the company's development teams using Cloudformation. The template will include building the VPC and networking components, installing a LAMP stack and securing the created resources. As per the AWS best practices what is the best way to design this template

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

Answer: D

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

NEW QUESTION 129
An enterprise wants to use a third-party SaaS application running on AWS.. The SaaS application needs to have access to issue several API commands to discover Amazon EC2 resources running within the enterprise's account. The enterprise has internal security policies that require any outside access to their environment must conform to the principles of least privilege and there must be controls in place to ensure that the credentials used by the SaaS vendor cannot be used by any other third party. Which of the following would meet all of these conditions?

A. From the AWS Management Console, navigate to the Security Credentials page and retrieve the access and secret key for your account.
B. Create an 1AM user within the enterprise account assign a user policy to the 1AM user that allows only the actions required by the SaaS applicatio
C. Create a new access and secret key for the user and provide these credentials to the SaaS provider.
D. Create an 1AM role for cross-account access allows the SaaS provider's account to assume the role and assign it a policy that allows only the actions required by the SaaS application.
E. Create an 1AM role for EC2 instances, assign it a policy that allows only the actions required tor the Saas application to work, provide the role ARN to the SaaS provider to use when launching their application instances.

Answer: C

Explanation:
Many SaaS platforms can access aws resources via a Cross account access created in aws. If you go to Roles in your identity management, you will see the ability to add a cross account role.

For more information on cross account role, please visit the below URL:
• http://docs.aws.amazon.com/IAM/latest/UserGuide/tuto rial_cross-account-with-roles.htm l

NEW QUESTION 130
Your company is getting ready to do a major public announcement of a social media site on AWS. The website is running on EC2 instances deployed across multiple Availability Zones with a Multi-AZ RDS MySQL Extra Large DB Instance. The site performs a high number of small reads and writes per second and relies on an eventual consistency model. After comprehensive tests you discover that there is read contention on RDS MySQL. Which are the best approaches to meet these requirements? Choose 2 answers from the options below

A. Deploy ElasticCache in-memory cache running in each availability zone
B. Implement sharding to distribute load to multiple RDS MySQL instances
C. Increase the RDS MySQL Instance size and Implement provisioned IOPS
D. Add an RDS MySQL read replica in each availability zone

Answer: AD

Explanation:
Implement Read Replicas and Clastic Cache
Amazon RDS Read Replicas provide enhanced performance and durability for database (DB) instances. This replication feature makes it easy to elastically scale out beyond the capacity constraints of a single DB Instance for read-heavy database workloads. You can create one or more replicas of a given source DB Instance and serve high-volume application read traffic from multiple copies of your data, thereby increasing aggregate read throughput.
For more information on Read Replica's, please visit the below link:
• https://aws.amazon.com/rds/details/read-replicas/
Amazon OastiCache is a web service that makes it easy to deploy, operate, and scale an in-memory data store or cache in the cloud. The service improves the performance of web applications by allowing you to retrieve information from fast, managed, in- memory data stores, instead of relying entirely on slower disk-based databases.
For more information on Amazon OastiCache, please visit the below link:
• https://aws.amazon.com/elasticache/

NEW QUESTION 135
Which of the following is incorrect when it comes to using the instances in an Opswork stack?

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

Answer: A

Explanation:
The AWS documentation mentions the following about Opswork stack
• A stack's instances can run either Linux or Windows.
A stack can have different Linux versions or distributions on different instances, but you cannot mix Linux and Windows instances.
• You can use custom AMIs (Amazon Machine Images), but they must be based on one of the AWS Ops Works Stacks-supported AMIs
• You can start and stop instances manually or have AWS OpsWorks Stacks automatically scale the number of instances. You can use time-based automatic scaling with any stack; Linux stacks also can use load-based scaling.
• In addition to using AWS OpsWorks Stacks to create Amazon EC2 instances, you can also register instances with a Linux stack that were created outside of AWS OpsWorks Stacks.
For more information on Opswork stacks, please visit the below link: http://docs.aws.amazon.com/opsworks/latest/userguide/workinginstances-os.html

NEW QUESTION 139
A user is trying to save some cost on the AWS services. Which of the below mentioned options will not help him save cost?

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

Answer: B

Explanation:
Option A is wrong because CBS volumes does have a costing aspect and hence deleting the volumes will save on cost

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

NEW QUESTION 144
Which of the following are components of the AWS Data Pipeline service. Choose 2 answers from the options given below

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

Answer: AB

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

NEW QUESTION 147
You work for a company that automatically tags photographs using artificial neural networks (ANNs), which run on GPUs using C++. You receive millions of images at a time, but only 3 times per day on average. These images are loaded into an AWS S3 bucket you control for you in a batch, and then the customer publishes a JSON-formatted manifest into another S3 bucket you control as well. Each image takes 10 milliseconds to process using a full GPU. Your neural network software requires 5 minutes to bootstrap. Image tags are JSON objects, and you must publish them to an S3 bucket. Which of these is the best system architectures for this system?

A. Create an OpsWorks Stack with two Layer
B. The first contains lifecycle scripts for launching and bootstrapping an HTTP API on G2 instances for image processing, and the second has an always-on instance which monitors the S3 manifest bucket for new file
C. When a new file is detected, request instances to boot on the artificial neural network laye
D. When the instances are booted and the HTTP APIs are up, submit processing requests to individual instances.
E. Make an S3 notification configuration which publishes to AWS Lambda on the manifest bucke
F. Make the Lambda create a CloudFormation Stack which contains the logic to construct an autoscaling worker tier of EC2 G2 instances with the artificial neural network code on each instanc
G. Create an SQS queue of the images in the manifes
H. Tear the stack down when the queue is empty.
I. Deploy your artificial neural network code to AWS Lambda as a bundled binary for the C++ extensio
J. Make an S3 notification configuration on the manifest, which publishes to another AWS Lambda running controller cod
K. This controller code publishes all the images in the manifest to AWS Kinesi
L. Your ANN code Lambda Function uses the Kinesis as an Event Sourc
M. The system automatically scales when the stream contains image events.
N. Create an Auto Scaling Load Balanced Elastic Beanstalk worker tier Application and Environmen
O. Deploy the artificial neural network code to G2 instances in this tie
P. Set the desired capacity to 1. Make the code periodically check S3 for new manifest
Q. When a new manifest is detected, push all of the images in the manifest into the SQS queue associated with the Elastic Beanstalk worker tier.

Answer: B

Explanation:
The best way to get notified when the images get sent to the bucket is the S3 Events. Here you don't need to provision infrastructure beforehand, and also since the source of S3 provides Event Handling, this should be used.
The AWS Documentation mentions the following on AWS S3 and Lambda
Amazon S3 can publish events (for example, when an object is created in a bucket) to AWS Lambda and invoke your Lambda function by passing the event data as a parameter. This integration enables
you to write Lambda functions that process Amazon S3 events. In Amazon S3, you add bucket notification configuration that identifies the type of event that you want Amazon S3 to publish and the Lambda function that you want to invoke.
For more information on AWS Lambda and S3 Events, please visit the link:
• http://docs.aws.amazon.com/lambda/latest/dg/with-s3, html

NEW QUESTION 149
A gaming company adopted AWS Cloud Formation to automate load-testing of theirgames. They have created an AWS Cloud Formation template for each gaming environment and one for the load- testing stack. The load-testing stack creates an Amazon Relational Database Service (RDS) Postgres database and two web servers running on Amazon Elastic Compute Cloud (EC2) that send HTTP requests, measure response times, and write the results into the database. A test run usually takes between 15 and 30 minutes. Once the tests are done, the AWS Cloud Formation stacks are torn down immediately. The test results written to the Amazon RDS database must remain accessible for visualization and analysis.
Select possible solutions that allow access to the test results after the AWS Cloud Formation load - testing stack is deleted.
Choose 2 answers.

A. Define an Amazon RDS Read-Replica in theload-testing AWS Cloud Formation stack and define a dependency relation betweenmaster and replica via the Depends On attribute.
B. Define an update policy to prevent deletionof the Amazon RDS database after the AWS Cloud Formation stack is deleted.
C. Define a deletion policy of type Retain forthe Amazon RDS resource to assure that the RDS database is not deleted with theAWS Cloud Formation stack.
D. Define a deletion policy of type Snapshotfor the Amazon RDS resource to assure that the RDS database can be restoredafter the AWS Cloud Formation stack is deleted.

E. Defineautomated backups with a backup retention period of 30 days for the Amazon RDSdatabase and perform point-in-time recovery of the database after the AWS CloudFormation stack is deleted.

Answer: CD

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

NEW QUESTION 154
Which of the following run command types are available for opswork stacks? Choose 3 answers from the options given below.

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

Answer: ABC

NEW QUESTION 158
Which of the following CLI commands is used to spin up new EC2 Instances?

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

Answer: A

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

NEW QUESTION 160
Your company has an e-commerce platform which is expanding all over the globe, you have EC2 instances deployed in multiple regions you want to monitor performance of all of these EC2 instances. How will you setup CloudWatch to monitor EC2 instances in multiple regions?

A. Createseparate dashboards in every region
B. Register!nstances running on different regions to CloudWatch
C. Haveone single dashboard to report metrics to CloudWatch from different region
D. Thisis not possible

Answer: C

Explanation:
You can monitor AWS resources in multiple regions using a single Cloud Watch dashboard. For example, you can create a dashboard that shows CPU utilization for an
CC2 instance located in the us-west-2 region with your billing metrics, which are located in the us- east-1 region.
For more information on Cloudwatch dashboard, please refer to the below url
http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/cross_region_dashboard.html

NEW QUESTION 161
A vendor needs access to your AWS account. They need to be able to read protected messages in a private S3 bucket. They have a separate AWS account. Which of the solutions below is the best way to do this?

A. Allowthe vendor to ssh into your EC2 instance and grant them an 1AM role with fullaccess to the bucket.
B. Createa cross-account 1AM role with permission to access the bucket, and grantpermission to use the role to the vendor AWS account.
C. Createan 1AM User with API Access Key
D. Give the vendor the AWS Access Key ID and AWSSecret Access Key for the user.
E. Createan S3 bucket policy that allows the vendor to read from the bucket from theirAWS account.

Answer: B

Explanation:

The AWS Documentation mentions the following on cross account roles
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 roles, please visit the below URL
http://docs.aws.amazon.com/IAM/latest/UserGuide/tuto rial_cross-account-with-roles.htm I https://docs.aws.amazon.com/AmazonS3/latest/dev/example-walkthroughs-managing-access- example2.html

NEW QUESTION 163
You are working for a company has an on-premise infrastructure. There is now a decision to move to AWS. The plan is to move the development environment first. There are a lot of custom based applications that need to be deployed for the development community. Which of the following can help to implement the application for the development team?
Choose 2 answers from the options below.

A. Create docker containers for the customapplication components.
B. Use OpsWorks to deploy the docker containers.
C. Use Elastic beanstalk to deploy the dockercontainers.
D. Use Cloudformation to deploy the dockercontainers.

Answer: AC

Explanation:
The AWS documentation states the following for docker containers on Elastic Beanstalk
Elastic Beanstalk supports the deployment of web applications from Docker containers. With Docker containers, you can define your own runtime environment. You can choose your own platform, programming language, and any application dependencies (such as package managers or tools), that aren't supported by other platforms. Docker containers are self-contained and include all the configuration information and software your web application requires to run.
For more information on docker containers and Elastic beanstalk, please visit the below URL
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker.html

NEW QUESTION 167
You are currently planning on using Autoscaling to launch instances which have an application installed. Which of the following methods will help ensure the instances are up and running in the shortest span of time to take in traffic from the users?

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

Answer: D

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

NEW QUESTION 169
Which of the following is not a component of Elastic Beanstalk?

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

Answer: C

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

NEW QUESTION 171
If you're trying to configure an AWS Elastic Beanstalk worker tier for easy debugging if there are problems finishing queue jobs, what should you configure?

A. ConfigureRolling Deployments.
B. ConfigureEnhanced Health Reporting.
C. ConfigureBlue-Green Deployments.
D. Configure a Dead Letter Queue.

Answer: D

Explanation:

The AWS documentation mentions the following on dead-letter queues
Amazon SQS supports dead-letter queues. A dead-letter queue is a queue that other (source) queues can target for messages that can't be processed (consumed)
successfully. You can set aside and isolate these messages in the dead-letter queue to determine why their processing doesn't succeed.
For more information on dead letter queues please visit the below link http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-dead-letter- queues. html

NEW QUESTION 175
You are Devops Engineer for a large organization. The company wants to start using Cloudformation templates to start building their resources in AWS. You are getting requirements for the templates from various departments, such as the networking, security, application etc. What is the best way to architect these Cloudformation templates.

A. Usea single Cloudformation template, since this would reduce the maintenanceoverhead on the templates itself.
B. Createseparate logical templates, for example, a separate template for networking,security, application et
C. Then nest the relevant templates.
D. Considerusing Elastic beanstalk to create your environments since Cloudformation is notbuilt for such customization.
E. Considerusing Opsworks to create your environments since Cloudformation is not builtfor such customization.

Answer: B

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

NEW QUESTION 180
You want to use Code Deploy to deploy code that is hosted on yourgithub repository. Which of the following additional services can help fulfil this requirement.

A. UsetheCodeP ipeline service
B. Usethe CodeCommit service
C. Usethe CodeBatch service
D. Usethe SQS service

Answer: A

Explanation:
The AWS Documentation mentions the following
AWS CodePipeline is a continuous delivery service you can use to model, visualize, and automate the steps required to release your software. You can quickly model and configure the different stages of a software release process. AWS CodePipeline automates the steps required to release your software changes continuously.
For more information on CodePipeline, please visit the below URL: http://docs.aws.a mazon.com/codepipeline/latest/userguide/concepts. htm l

NEW QUESTION 185
Which of the following Cloudformation helper scripts can help install packages on EC2 resources

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

Answer: A

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

NEW QUESTION 189
Which of the following CLI commands can be used to describe the stack resources.

A. awscloudformationdescribe-stack
B. awscloudformationdescribe-stack-resources
C. awscloudformation list-stack-resources
D. awscloudformation list-stack

Answer: C

Explanation:
Answer - C

This is given in the AWS Documentation list-stack-resources
Description
Returns descriptions of all resources of the specified stack.
For deleted stacks, ListStackResources returns resource information for up to 90 days after the stack has been deleted.
See also: AWS API Documentation
See 'aws help' for descriptions of global parameters.
list-stack-resources is a paginated operation. Multiple API calls may be issued in order to retrieve the entire data set of results. You can disable pagination by providing the —no-paginate argument. When using —output text and the —query argument on a paginated response, the —query argument must extract data from the results of the following query expressions: StackResourceSummaries For more information on the CLI command, please visit the below URL:
http://docs.aws.amazon.com/cli/latest/reference/cloudformation/list-stack-resources.html

NEW QUESTION 194
You have just developed a new mobile application that handles analytics workloads on large scale datasets that are stored on Amazon Redshift. Consequently, the application needs to access Amazon Redshift tables. Which of the below methods would be the best, both practically and security-wise, to access the tables?
Choose the correct answer from the options below

A. Createan 1AM user and generate encryption keys for that use
B. Create a policy for RedShiftread- only acces
C. Embed the keys in the application.
D. Createa HSM client certificate in Redshift and authenticate using this certificate.
E. Createa RedShift read-only access policy in 1AM and embed those credentials in theapplication.
F. Useroles that allow a web identity federated user to assume a role that allowsaccess to the RedShift table by providing temporary credentials.

Answer: D

Explanation:
For access to any AWS service, the ideal approach for any application is to use Roles. This is the first preference. Hence option A and C are wrong.
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.
"When you write such an app, you'll make requests to AWS services that must be signed with an AWS access key. However, we strongly recommend that you do not 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 196
What are the benefits when you implement a Blue Green deployment for your infrastructure or application level changes. Choose 3 answers from the options given below

A. Nearzero-downtime release for new changes
B. Betterrollback capabilities
C. Abilityto deploy with higher risk
D. Goodturnaround time for application deployments

Answer: ABD

Explanation:
The AWS Documentation mentions the following
Blue/green deployments provide near zero-downtime release and rollback capabilities. The fundamental idea behind blue/green deployment is to shift traffic between two identical environments that are running different versions of your application. The blue environment represents the current application version serving production traffic. In parallel, the green environment is staged running a different version of your application. After the green environment is ready and tested, production traffic is redirected from blue to green.
For more information on Blue Green deployments please see the below link:
• https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 199
As part of your continuous deployment process, your application undergoes an I/O load performance test before it is deployed to production using new AMIs. The application uses one Amazon EBS PIOPS volume per instance and requires consistent I/O performance.
Which of the following must be carried out to ensure that I/O load performance tests yield the correct results in a repeatable manner?

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

Answer: B

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

NEW QUESTION 203
Your current log analysis application takes more than four hours to generate a report of the top 10 users of your web application. You have been asked to implement a system that can report this information in real time, ensure that the report is always up to date, and handle increases in the number of requests to your web application. Choose the option that is cost-effective and can fulfill the
requirements.

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

Answer: C

Explanation:
The AWS Documentation mentions the below
Amazon Kinesis makes it easy to collect, process, and analyze real-time, streaming data so you can get timely insights and react quickly to new information. Amazon
Kinesis offers key capabilities to cost effectively process streaming data at any scale, along with the flexibility to choose the tools that best suit the requirements of your application. With Amazon Kinesis, you can ingest real-time data such as application logs, website clickstreams, IoT telemetry data, and more into your databases, data lakes and data warehouses, or build your own real-time applications using this data.
Amazon Kinesis enables you to process and analyze data as it
arrives and respond in real-time instead of having to wait until all your data is collected before the processing can begin.
For more information on AWS Kinesis please see the below link:
• https://aws.amazon.com/kinesis/

NEW QUESTION 207
You are managing an application that contains Go as the front end, MongoDB for document management and is hosted on a relevant Web server. You pre-bake AMI'S with the latest version of the Web server, then user the User Data section to setup the application. You now have a change to the underlying Operating system version and need to deploy that accordingly. How can this be done
in the easiest way possible.

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

Answer: C

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

NEW QUESTION 212
You are managing the development of an application that uses DynamoDB to store JSON data. You have already set the Read and Write capacity of the DynamoDB table. You are unsure of the amount of the traffic that will be received by the application during the deployment time. How can you ensure that the DynamoDB is not highly throttled and does not become a bottleneck for the application? Choose 2 answers from the options below.

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

Answer: AC

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

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

For more information on monitoring DynamoDB please seethe below link:
• http://docs.aws.a mazon.com/amazondynamodb/latest/developerguide/monitori ng- cloudwatch.html

NEW QUESTION 216
When deploying applications to Elastic Beanstalk, which of the following statements is false with regards to application deployment

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

Answer: B

Explanation:
The AWS Documentation mentions
When you use the AWS Clastic Beanstalk console to deploy a new application or an application version, you'll need to upload a source bundle. Your source bundle must meet the following requirements:
Consist of a single ZIP file or WAR file (you can include multiple WAR files inside your ZIP file) Not exceed 512 MB
Not include a parent folder or top-level directory (subdirectories are fine)
For more information on deploying applications to Clastic Beanstalk please see the below link: http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/applications-sourcebundle.html

NEW QUESTION 221
You have an Autoscaling Group which is launching a set of t2.small instances. You now need to replace those instances with a larger instance type. How would you go about making this change in an ideal manner?

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

Answer: C

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

NEW QUESTION 222
Your application has an Auto Scaling group of three EC2 instances behind an Elastic Load Balancer. Your Auto Scalinggroup was updated with a new launch configuration that refers to an updated AMI. During the deployment, customers complained that they were receiving several errors even though all instances passed the ELB health checks. How can you prevent this from happening again?

A. Createa new ELB and attach the Autoscaling Group to the ELB
B. Createa new launch configuration with the updated AMI and associate it with the AutoScaling grou
C. Increase the size of the group to six and when instances becomehealthy revert to three.
D. Manuallyterminate the instances with the older launch configuration.
E. Updatethe launch configuration instead of updating the Autoscaling Group

Answer: B

Explanation:
An Auto Scaling group is associated with one launch configuration at a time, and you can't modify a launch configuration after you've created it. To change the launch configuration for an Auto Scaling group, you can use an existing launch configuration as the basis for a new launch configuration and then update the Auto Scaling group to use the new launch configuration.
After you change the launch configuration for an Auto Scaling group, any new instances are launched using the new configuration options, but existing instances are not affected.
Then to ensure the new instances are launches, change the size of the Autoscaling Group to 6 and once the new instances are launched, change it back to 3.
For more information on instances scale-in process and Auto Scaling Group's termination policies please view the following link:
• https://docs^ws.amazon.com/autoscaling/ec2/userguide/as-instance-termination.html#default- termination-policy For more information on changing the launch configuration please see the below link:
• http://docs.aws.amazon.com/autoscaling/latest/userguide/change-launch-config.html

NEW QUESTION 227
Which of the following is a container for metrics in Cloudwatch?

A. MetricCollection
B. Namespaces
C. Packages
D. Locale

Answer: B

Explanation:
The AWS Documentation mentions the following
Cloud Watch namespaces are containers for metrics. Metrics in different namespaces are isolated from each other, so that metrics from different applications are not mistakenly aggregated into the same statistics. All AWS services that provide Amazon Cloud Watch data use a namespace string, beginning with "AWS/". When
you create custom metrics, you must also specify a namespace as a container for custom metrics. The following services push metric data points to Cloud Watch.
For more information on Cloudwatch namespaces, please visit the below URL: ttp://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/aws-namespaces.html

NEW QUESTION 230
By default in Opswork, how many application versions can you rollback up to?

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

Answer: D

Explanation:

The AWS Documentation mentions the following Restores the previously deployed app version. For example, if you have deployed the app three times and then run Rollback, the server will serve the app from the second deployment. If you run Rollback again, the server will serve the app from the first deployment. By default, AWS OpsWorks Stacks stores the five most recent deployments, which allows you to roll back up to four versions. If you exceed the number of stored versions, the command fails and leaves the oldest version in place.
For more information on Opswork app deployment, please visit the below U RL: http://docs.aws.amazon.com/opsworks/latest/userguide/workingapps-deploying.html

NEW QUESTION 233
You have defined a Linux based instance stack in Opswork. You now want to attach a database to the Opswork stack. Which of the below is an important step to ensure that the application on the Linux instances can communicate with the database

A. Addanother stack with the database layer and attach it to the application stack.
B. ConfigureSSL so that the instance can communicate with the database
C. Addthe appropriate driver packages to ensure the application can work with thedatabase
D. Configuredatabase tags for the Opswork application layerOpswork application layer

Answer: C

Explanation:
The AWS documentation mentions the below point Important
For Linux stacks, if you want to associate an Amazon RDS service layer with your app, you must add the appropriate driver package to the associated app server layer,
as follows:
1. Click Layers in the navigation pane and open the app server's Recipes tab.
2. Click Edit and add the appropriate driver package to OS Packages. For example, you should specify mysql if the layer contains Amazon Linux instances and mysql-client if the layer contains Ubuntu instances.
3. Save the changes and redeploy the app.
For more information on Opswork app connectivity, please visit the below URL: http://docs.aws.amazon.com/opsworks/latest/userguide/workingapps-connectdb.html

NEW QUESTION 235
.....

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