Exam Questions DOP-C01

AWS Certified DevOps Engineer- Professional

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

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

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

Answer: D

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

NEW QUESTION 2
You have deployed an application to AWS which makes use of Autoscaling to launch new instances. You now want to change the instance type for the new instances. Which of the following is one of the action items to achieve this deployment?

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

Answer: C

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

NEW QUESTION 3
Your application stores sensitive information on an EBS volume attached to your EC2 instance. How can you protect your information? Choose two answers from the options given below

A. Unmount the EBS volume, take a snapshot and encrypt the snapsho
B. Re-mount the Amazon EBS volume
C. It is not possible to encrypt an EBS volume, you must use a lifecycle policy to transfer data to S3 forencryption.
D. Copy the unencrypted snapshot and check the box to encrypt the new snapsho
E. Volumes restored from this encrypted snapshot will also be encrypted.
F. Create and mount a new, encrypted Amazon EBS volum
G. Move the data to the new volum
H. Delete the old Amazon EBS volume *t

Answer: CD

Explanation:
These steps are given in the AWS documentation
To migrate data between encrypted and unencrypted volumes
1) Create your destination volume (encrypted or unencrypted, depending on your need).
2) Attach the destination volume to the instance that hosts the data to migrate.
3) Make the destination volume available by following the procedures in Making an Amazon EBS Volume Available for Use. For Linux instances, you can create a mount point at /mnt/destination and mount the destination volume there.
4) Copy the data from your source directory to the destination volume. It may be most convenient to use a bulk-copy utility for this.
To encrypt a volume's data by means of snapshot copying
1) Create a snapshot of your unencrypted CBS volume. This snapshot is also unencrypted.
2) Copy the snapshot while applying encryption parameters. The resulting target snapshot is encrypted.
3) Restore the encrypted snapshot to a new volume, which is also encrypted.
For more information on EBS Encryption, please refer to the below document link: from AWS
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html

NEW QUESTION 4
Your company has multiple applications running on AWS. Your company wants to develop a tool that notifies on-call teams immediately via email when an alarm is triggered in your environment. You have multiple on-call teams that work different shifts, and the tool should handle notifying the correct teams at the correct times. How should you implement this solution?

A. Create an Amazon SNS topic and an Amazon SQS queue
B. Configure the Amazon SQS queue as a subscriber to the Amazon SNS topic.Configure CloudWatch alarms to notify this topic when an alarm is triggere

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

Answer: D

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

NEW QUESTION 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/WhatI 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 have an application consisting of a stateless web server tier running on Amazon EC2 instances behind load balancer, and are using Amazon RDS with read replicas. Which of the following methods should you use to implement a self-healing and cost-effective architecture? Choose 2 answers from the optionsgiven

below

A. Set up a third-party monitoring solution on a cluster of Amazon EC2 instances in order to emit custom Cloud Watch metrics to trigger the termination of unhealthy Amazon EC2 instances.
B. Set up scripts on each Amazon EC2 instance to frequently send ICMP pings to the load balancer in order to determine which instance is unhealthy and replace it.
C. Set up an Auto Scalinggroup for the web server tier along with an Auto Scaling policy that uses the Amazon RDS DB CPU utilization Cloud Watch metric to scale the instances.
D. Set up an Auto Scalinggroup for the web server tier along with an Auto Scaling policy that uses the Amazon EC2 CPU utilization CloudWatch metric to scale the instances.
E. Use a larger Amazon EC2 instance type for the web server tier and a larger DB instance type for the data storage layer to ensure that they don't become unhealthy.
F. Set up an Auto Scalinggroup for the database tier along with an Auto Scaling policy that uses the Amazon RDS read replica lag CloudWatch metric to scale out the Amazon RDS read replicas.
G. Use an Amazon RDS Multi-AZ deployment.

Answer: DG

Explanation:
The scaling of CC2 Instances in the Autoscaling group is normally done with the metric of the CPU utilization of the current instances in the Autoscaling group
For more information on scaling in your Autoscaling Group, please refer to the below link:
• http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scaling-simple-step.html
Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a Multi- AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Cach AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure, Amazon RDS performs an automatic failover to the standby (or to a read replica in the case of Amazon Aurora), so that you can resume database operations as soon as the failover is complete. For more information on RDS Multi-AZ please refer to the below link:
https://aws.amazon.com/rds/details/multi-az/
Option A is invalid because if you already have in-built metrics from Cloudwatch, why would you want to spend more in using a a third-party monitoring solution.
Option B is invalid because health checks are already a feature of AWS CLB
Option C is invalid because the database CPU usage should not be used to scale the web tier.
Option C is invalid because increasing the instance size does not always guarantee that the solution will not become unhealthy.
Option F is invalid because increasing Read-Replica's will not suffice for write operations if the primary DB fails.

NEW QUESTION 8
You have the following application to be setup in AWS
1) A web tier hosted on EC2 Instances
2) Session data to be written to DynamoDB
3) Log files to be written to Microsoft SQL Server
How can you allow an application to write data to a DynamoDB table?

A. Add an 1AM user to a running EC2 instance.
B. Add an 1AM user that allows write access to the DynamoDB table.
C. Create an 1AM role that allows read access to the DynamoDB table.
D. Create an 1AM role that allows write access to the DynamoDB table.

Answer: D

Explanation:
I AM roles are designed so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that
the applications use. Instead of creating and distributing your AWS credentials For more information on 1AM Roles please refer to the below link:
http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/iam-roles-for-amazon-ec2.html

NEW QUESTION 9
You are a DevOps engineer for a company. You have been requested to create a rolling deployment solution that is cost-effective with minimal downtime. How should you achieve this? Choose two answers from the options below

A. Re-deploy your application using a CloudFormation template to deploy Elastic Beanstalk
B. Re-deploy with a CloudFormation template, define update policies on Auto Scalinggroups in your CloudFormation template
C. Use UpdatePolicy attribute to specify how CloudFormation handles updates to Auto Scaling Group resource.
D. After each stack is deployed, tear down the old stack

Answer: BC

Explanation:
The AWS::AutoScaling::AutoScalingGroup resource supports an UpdatePolicy attribute. This is used to define how an Auto Scalinggroup resource is updated when
an update to the Cloud Formation stack occurs. A common approach to updating an Auto Scaling group is to perform a rolling update, which is done by specifying the
AutoScalingRollingUpdate policy. This retains the same Auto Scalinggroup and replaces old instances with new ones, according to the parameters specified.
Option A is invalid because it is not efficient to use Cloudformation to use Clastic Beanstalk.
Option D is invalid because this is an inefficient process to tear down stacks when there are stack policies available
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 10
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 10
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 12
You have an Auto Scaling group with 2 AZs. One AZ has 4 EC2 instances and the other has 3 EC2 instances. None of the instances are protected from scale in. Based on the default Auto Scaling termination policy what will happen?

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

Answer: D

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

NEW QUESTION 13
You are doing a load testing exercise on your application hosted on AWS. While testing your Amazon RDS MySQL DB instance, you notice that when you hit 100% CPU utilization on it, your application becomes non- responsive. Your application is read-heavy. What are methods to scale your data tier to meet the application's needs? Choose three answers from the options given below

A. Add Amazon RDS DB read replicas, and have your application direct read queries to them.
B. Add your Amazon RDS DB instance to an Auto Scalinggroup and configure your Cloud Watch metricbased on CPU utilization.
C. Use an Amazon SQS queue to throttle data going to the Amazon RDS DB instance.
D. Use ElastiCache in front of your Amazon RDS DB to cache common queries.
E. Shard your data set among multiple Amazon RDS DB instances.
F. Enable Multi-AZ for your Amazon RDS DB instance.

Answer: ADE

Explanation:
Amazon RDS Read Replicas provide enhanced performance and durability for database (DB) instances. This replication feature makes it easy to elastically scale out
beyond the capacity constraints of a single DB Instance for read-heavy database workloads. You can create one or more replicas of a given source DB Instance and
serve high-volume application read traffic from multiple copies of your data, thereby increasing aggregate read throughput
For more information on Read Replica's please refer to the below link:
? https://aws.amazon.com/rds/details/read-replicas/
Sharding is a common concept to split data across multiple tables in a database For more information on sharding please refer to the below link:
https://forums.awsa mazon.com/thread jspa?messageID=203052

Amazon OastiCache is a web service that makes it easy to deploy, operate, and scale an in-memory data store or cache in the cloud. The service improves the performance of web applications by allowing you to retrieve information from fast, managed, in-memory data stores, instead of relying entirely on slower disk-based databases

Amazon OastiCache is an in-memory key/value store that sits between ycbetappiicipiJGra arcdalie data store (database) that it accesses. Whenever your application requests data, it first makes the request to the DastiCache cache. If the data exists in the cache and is current, OastiCache returns the data to your application. If the data does not exist in the cache, or the data in the cache has expired, your application requests the data from your data store which returns the data to your application. Your application then writes the data received from the store to the cache so it can be more quickly retrieved next time it is requested. For more information on Elastic Cache please refer to the below link:

https://aws.amazon.com/elasticache/

Option B is not an ideal way to scale a database

Option C is not ideal to store the data which would go into a database because of the message size Option F is invalid because Multi-AZ feature is only a failover option

NEW QUESTION 17
You have an Auto Scaling group with an Elastic Load Balancer. You decide to suspend the Auto Scaling AddToLoadBalancer for a short period of time. What will happen to the instances launched during the suspension period?

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

Answer: C

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

NEW QUESTION 19
You have a current Clouformation template defines in AWS. You need to change the current alarm threshold defined in the Cloudwatch alarm. How can you achieve this?

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

Answer: C

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

NEW QUESTION 24
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/WhatIsCloudWatchLogs.html

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

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

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

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

NEW QUESTION 27
You have a large number of web servers in an Auto Scalinggroup behind a load balancer. On an hourly basis, you want to filter and process the logs to collect data on unique visitors, and then put that data in a durable data store in order to run reports. Web servers in the Auto Scalinggroup are constantly launching and terminating based on your scaling policies, but you do not want to lose any of the log data from these servers during a stop/termination initiated by a user or by Auto Scaling. What two approaches will meet these requirements? Choose two answers from the optionsgiven below.

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

Answer: AC

Explanation:
You can use the Cloud Watch Logs agent installer on an existing CC2 instance to install and configure the Cloud Watch Logs agent.
For more information, please visit the below link:
• http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/Qu ickStartCC2lnstance.html
You can publish your own metrics to Cloud Watch using the AWS CLI or an API. For more information, please visit the below link:
• http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.htmlI Amazon Redshift is a fast, fully managed data warehouse that makes it simple and cost-effective to analyze all your data using standard SQL and your existing Business Intelligence (BI) tools. It allows you to run complex analytic queries against petabytes of structured data, using sophisticated query optimization, columnar storage on high-performance local disks, and massively parallel query execution. Most results come back in seconds. For more information on copying data from S3 to redshift, please refer to the below link:
• http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/dp-copydata- redshift html

NEW QUESTION 32
You have a web application that's developed in Node.js The code is hosted in Git repository. You want to now deploy this application to AWS. Which of the below 2 options can fulfil this requirement.

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

Answer: CD

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

NEW QUESTION 36
You have an application running on Amazon EC2 in an Auto Scaling group. Instances are being bootstrapped dynamically, and the bootstrapping takes over 15 minutes to complete. You find that instances are reported by Auto Scaling as being In Service before bootstrapping has completed. You are receiving application alarms related to new instances before they have completed bootstrapping, which is causing confusion. You find the cause: your application monitoring tool is

polling the Auto Scaling Service API for instances that are In Service, and creating alarms for new previously unknown instances. Which of the following will ensure that new instances are not added to your application monitoring tool before bootstrapping is completed?

A. Create an Auto Scaling group lifecycle hook to hold the instance in a pending: wait state until your bootstrapping is complet
B. Once bootstrapping is complete, notify Auto Scaling to complete the lifecycle hook and move the instance into a pending:proceed state.
C. Use the default Amazon Cloud Watch application metrics to monitor your application's heal t
D. Configure an Amazon SNS topic to send these Cloud Watch alarms to the correct recipients.
E. Tag all instances on launch to identify that they are in a pending stat
F. Change your application monitoring tool to look for this tag before adding new instances, and the use the Amazon API to set the instance state to 'pending' until bootstrapping is complete.
G. Increase the desired number of instances in your Auto Scaling group configuration to reduce the time it takes to bootstrap future instances.

Answer: A

Explanation:
Auto Scaling lifecycle hooks enable you to perform custom actions as Auto Scaling launches or terminates instances. For example, you could install or configure software on newly launched instances, or download log files from an instance before it terminates. After you add lifecycle hooks to your Auto Scaling group, they work as follows:
1. Auto Scaling responds to scale out events by launching instances and scale in events by terminating instances.
2. Auto Scaling puts the instance into a wait state (Pending:Wait orTerminating:Wait). The instance remains in this state until either you tell Auto Scaling to continue or the timeout period ends.
For more information on rolling updates, please visit the below link:
• http://docs.aws.amazon.com/autoscaling/latest/userguide/lifecycle-hooks.htm l

NEW QUESTION 40
You currently have an Auto Scaling group with an Elastic Load Balancer and need to phase out all instances and replace with a new instance type. What are 2 ways in which this can be achieved.
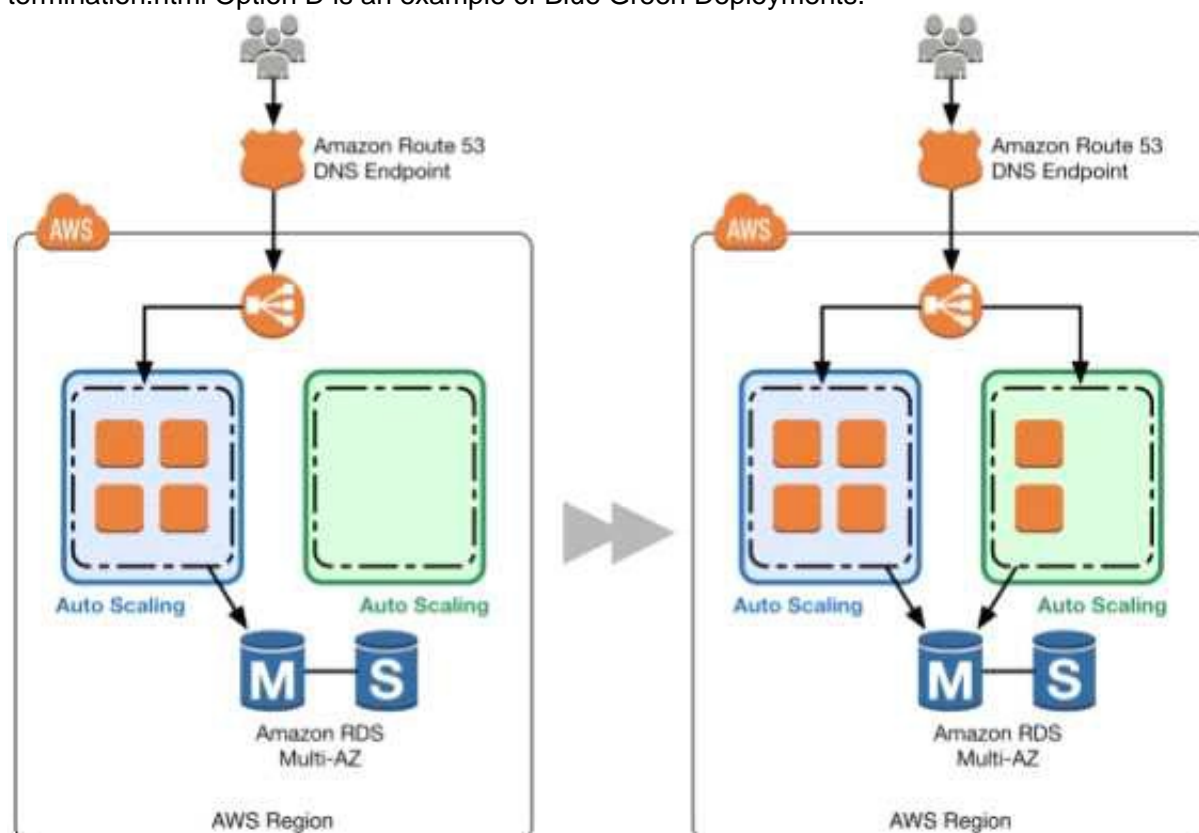
A. Use Newest In stance to phase out all instances that use the previous configuration.
B. Attach an additional ELB to your Auto Scaling configuration and phase in newer instances while removing older instances.
C. Use OldestLaunchConfiguration to phase out all instances that use the previous configuratio
D. V
E. Attach an additional Auto Scaling configuration behind the ELB and phase in newer instances while removing older instances.

Answer: CD

Explanation:
When using the OldestLaunchConfiguration policy Auto Scaling terminates instances that have the oldest launch configuration. This policy is useful when you're updating a group and phasing out the instances from a previous configuration.
For more information on Autoscaling instance termination, please visit the below URL: http://docs.aws.amazon.com/autoscaling/latest/userguide/as-instance-termination.html Option D is an example of Blue Green Deployments.

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 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 44
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 47
When thinking of AWS Elastic Beanstalk's model, which is true?

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

Answer: C

Explanation:
The first step in using Elastic Beanstalk is to create an application, which represents your web application in AWS. In Elastic Beanstalk an application serves as a container for the environments that run your web app, and versions of your web app's source code, saved configurations, logs and other artifacts that you create while using Elastic Beanstalk.
For more information on Applications, please refer to the below link: http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/applications.html
Deploying a new version of your application to an environment is typically a fairly quick process. The new source bundle is deployed to an instance and extracted, and the the web container or application server picks up the new version and restarts if necessary. During deployment, your application might still become unavailable to users for a few seconds. You can prevent this by configuring your environment to use rolling deployments to deploy the new version to instances in batches. For more information on deployment, please refer to the below link: http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.de ploy-existing-version, html

NEW QUESTION 49
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 50
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/
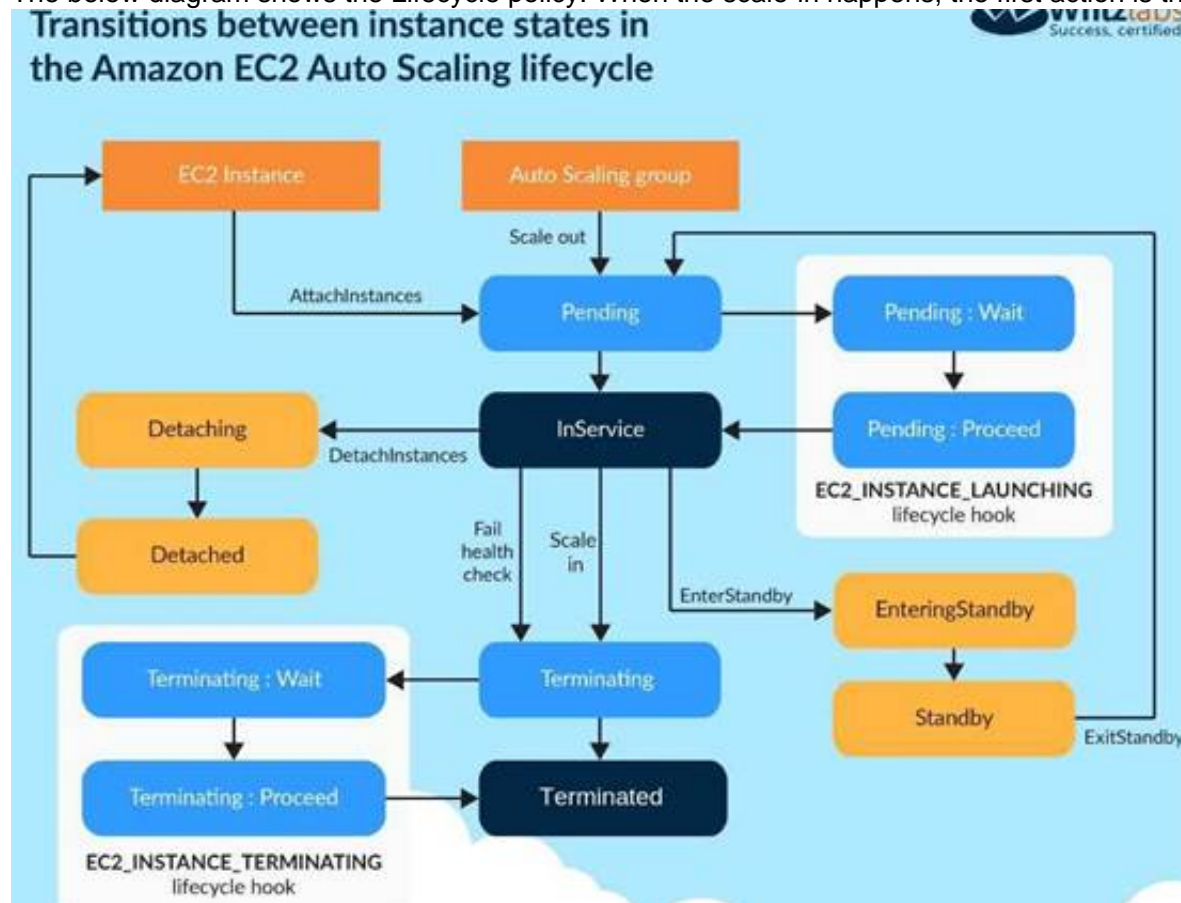
NEW QUESTION 53
For AWS Auto Scaling, what is the first transition state an instance enters after leaving steady state when scaling in due to health check failure or decreased load?

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

Answer: A

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

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

NEW QUESTION 58
You are using Chef in your data center. Which service is designed to let the customer leverage existing Chef recipes in AWS?

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

Answer: B

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

NEW QUESTION 61
Your company releases new features with high frequency while demanding high application availability. As part of the application's A/B testing, logs from each updated Amazon EC2 instance of the application need to be analyzed in near real-time, to ensure that the application is working
flawlessly after each deployment. If the logs show any anomalous behavior, then the application version of the instance is changed to a more stable one. Which of the following methods should you use for shipping and analyzing the logs in a highly available manner?

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

Answer: C

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

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

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

Answer: A

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

NEW QUESTION 67
Which of these is not an instrinsic function in AWS CloudFormation?

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

Answer: D

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

NEW QUESTION 69
You have been given a business requirement to retain log files for your application for 10 years. You need to regularly retrieve the most recent logs for troubleshooting. Your logging system must be cost-effective, given the large volume of logs. What technique should you use to meet these requirements?

A. Store your log in Amazon CloudWatch Logs.
B. Store your logs in Amazon Glacier.
C. Store your logs in Amazon S3, and use lifecycle policies to archive to Amazon Glacier.
D. Store your logs on Amazon EBS, and use Amazon EBS snapshots to archive them.

Answer: C

Explanation:
Option A is invalid, because cloud watch will not store the logs indefinitely and secondly it won't be the cost effective option.
Option B is invalid, because it won't server the purpose of regularly retrieve the most recent logs for troubleshooting. You will need to pay more to retrieve the logs faster from this storage.
Option D is invalid, because it is not an ideal or cost effective option.
You can define lifecycle configuration rules for objects that have a well-defined lifecycle. For example: if you are uploading periodic logs to your bucket, your application might need these logs for a week or a month after creation, and after that you might want to delete them.
Some documents are frequently accessed for a limited period of time. After that, these documents are less frequently accessed. Over time, you might not need real-time access to these objects, but your organization or regulations might require you to archive them for a longer period and then optionally delete them later.
You might also upload some types of data to Amazon S3 primarily for archival purposes, for example digital media archives, financial and healthcare records, raw genomics sequence data, long-term database backups, and data that must be retained for regulatory compliance.
For more information on Lifecycle management please refer to the below link: http://docs.aws.amazon.com/AmazonS3/latest/dev/object-lifecycle-mgmt.htmI
Note:
Option C is the cheapest option, but Cloud watch can store logs indefinetly or between 10 years and
one day.
"Log Retention—By default, logs are kept indefinitely and never expire. You can adjust the retention policy for each log group, keeping the indefinite retention, or choosing a retention periods between 10 years and one day." https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html

NEW QUESTION 70
There is a requirement to monitor API calls against your AWS account by different users and entities. There needs to be a history of those calls. The history of those calls are needed in in bulk for later review. Which 2 services can be used in this scenario

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

Answer: C

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

NEW QUESTION 74
You want to pass queue messages that are 1GB each. How should you achieve this?

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

Answer: B

Explanation:
You can manage Amazon SQS messages with Amazon S3. This is especially useful for storing and consuming messages with a message size of up to 2 GB. To manage
Amazon SQS messages with Amazon S3, use the Amazon SQS Extended Client Library for Java. Specifically, you use this library to:
• Specify whether messages are always stored in Amazon S3 or only when a message's size exceeds 256 KB.
• Send a message that references a single message object stored in an Amazon S3 bucket.
• Get the corresponding message object from an Amazon S3 bucket.
• Delete the corresponding message object from an Amazon S3 bucket.
For more information on processing large messages for SQS, please visit the below URL:
http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-s3- messages. html
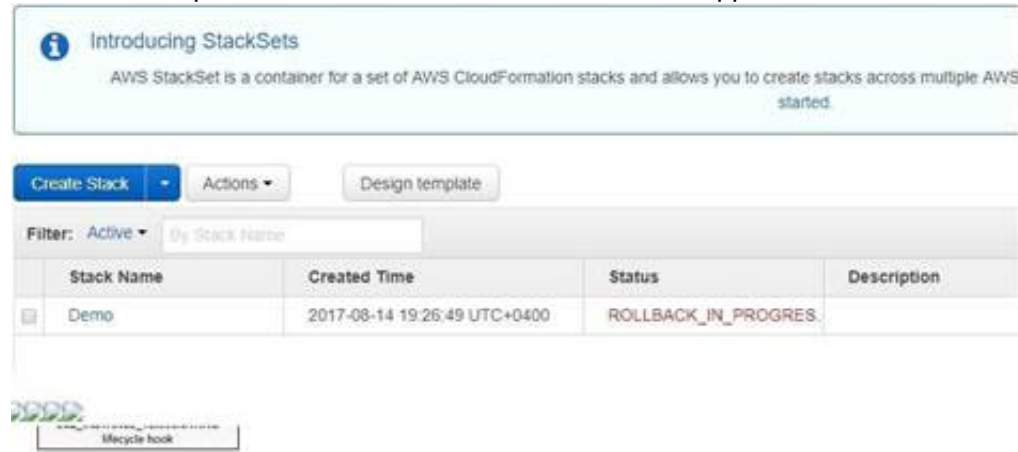
NEW QUESTION 77
You have deployed a Cloudformation template which is used to spin up resources in your account. Which of the following status in Cloudformation represents a failure.

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

Answer: C

Explanation:
AWS Cloud Formation provisions and configures resources by making calls to the AWS services that are described in your template.
After all the resources have been created, AWS Cloud Formation reports that your stack has been created. You can then start using the resources in your stack. If stack creation fails, AWS CloudFormation rolls back your changes by deleting the resources that it created.
The below snapshot from Cloudformation shows what happens when there is an error in the stack creation.

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

NEW QUESTION 80
You are building out a layer in a software stack on AWS that needs to be able to scale out to react to increased demand as fast as possible. You are running the code on EC2 instances in an Auto Scaling Group behind an ELB. Which application code deployment method should you use?

A. SSH into new instances that come online, and deploy new code onto the system by pulling it from an S3 bucket, which is populated by code that you refresh from source control on new pushes.
B. Bake an AMI when deploying new versions of code, and use that AMI for the Auto Scaling Launch Configuration.
C. Create a Dockerfile when preparing to deploy a new version to production and publish it to S3. Use UserData in the Auto Scaling Launch configuration to pull down the Dockerfile from S3 and run it when new instances launch.
D. Create a new Auto Scaling Launch Configuration with UserData scripts configured to pull the latest code at all times.

Answer: B

Explanation:
Since the time required to spin up an instance is required to be fast, its better to create an AMI rather than use User Data. When you use User Data, the script will be
run during boot up, and hence this will be slower.
An Amazon Machine Image (AMI) provides the information required to launch an instance, which is a virtual server in the cloud. You specify an AM I when you launch
an instance, and you can launch as many instances from the AMI as you need. You can also launch instances from as many different AMIs as you need.
For more information on the AMI, please refer to the below link:
• http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/AMIs.html

NEW QUESTION 81
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 83
You run accounting software in the AWS cloud. This software needs to be online continuously during the day every day of the week, and has a very static requirement for compute resources. You also have other, unrelated batch jobs that need to run once per day at anytime of your choosing. How should you minimize cost?

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

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

Answer: A

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

NEW QUESTION 87
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 90
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 94
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 96
You need to run a very large batch data processingjob one time per day. The source data exists
entirely in S3, and the output of the processingjob should also be written to S3 when finished. If you need to version control this processingjob and all setup and
teardown logic for the system, what approach should you use?.

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

Answer: B

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

NEW QUESTION 100
Which of the following tools does not directly support AWS OpsWorks, for monitoring your stacks?

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

Answer: A

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

NEW QUESTION 103
There is a very serious outage at AWS. EC2 is not affected, but your EC2 instance deployment scripts
stopped working in the region with the outage. What might be the issue?

A. The AWS Console is down, so your CLI commands do not work.
B. S3 is unavailable, so you can't create EBS volumes from a snapshot you use to deploy new volumes.
C. AWS turns off the DeployCode API call when there are major outages, to protect from system floods.
D. None of the other answers make sens
E. If EC2 is not affected, it must be some other issue.

Answer: B

Explanation:
The CBS Snapshots are stored in S3, so if you have an scripts which deploy CC2 Instances, the CBS volumes need to be constructed from snapshots stored in S3.
You can back up the data on your Amazon CBS volumes to Amazon S3 by taking point-in-time snapshots. Snapshots are incremental backups, which means that only the blocks on the device that have changed after your most recent snapshot are saved. This minimizes the time required to create the snapshot and saves on storage costs by not duplicating data. When you delete a snapshot, only the data unique to that snapshot is removed. Cach snapshot contains all of the information needed to restore your data (from the moment when the snapshot was taken) to a new CBS volume. For more information on CBS Snapshots, please visit the below URL:
• http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/ CBSSnapshots.htm I

NEW QUESTION 107
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 110
You have an asynchronous processing application usingan Auto Scaling Group and an SQS Queue. The Auto Scaling Group scales according to the depth of the job queue. The completion velocity of the jobs has gone down, the Auto Scaling Group size has maxec out, but the inbound job velocity did not increase. What is a possible issue?

A. Some of the new jobs coming in are malformed and unprocessable.
B. The routing tables changed and none of the workers can process events anymore.
C. Someone changed the 1AM Role Policy on the instances in the worker group and broke permissions to access the queue.
D. The scaling metric is not functioning correctly.

Answer: A

Explanation:
This question is more on the grounds of validating each option
Option B is invalid, because the Route table would have an effect on all worker processes and no jobs would have been completed.
Option C is invalid because if the 1AM Role was invalid then no jobs would be completed.
Option D is invalid because the scaling is happening, its just that the jobs are not getting completed. For more information on Scaling on Demand, please visit the below URL:
• http://docs.aws.a mazon.com/autoscaling/latest/userguide/as-scale-based-on-demand.html

NEW QUESTION 112
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 116
Your team wants to begin practicing continuous delivery using CloudFormation, to enable automated builds and deploys of whole, versioned stacks or stack layers. You have a 3-tier, mission-critical system. Which of the following is NOT a best practice for using CloudFormation in a continuous delivery environment?

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

Answer: B

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

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

NEW QUESTION 119
You need to deploy an AWS stack in a repeatable manner across multiple environments. You have selected CloudFormation as the right tool to accomplish this, but have found that there is a resource type you need to create and model, but is unsupported by CloudFormation. How should you overcome this challenge?

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

Answer: D

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

NEW QUESTION 122
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 127
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 128
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 132
Your development team is using access keys to develop an application that has access to S3 and DynamoDB. A new security policy has outlined that the credentials should not be older than 2 months, and should be rotated. How can you achieve this

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

Answer: B

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

NEW QUESTION 134
You currently have an application deployed via Elastic Beanstalk. You are now deploying a new application and have ensured that Elastic beanstalk has detached the current instances and deployed and reattached new instances. But the new instances are still not receiving any sort of traffic. Why is this the case.

A. The instances are of the wrong AMI, hence they are not being detected by the ELB.
B. It takes time for the ELB to register the instances, hence there is a small timeframe before your instances can start receiving traffic
C. You need to create a new Elastic Beanstalk application, because you cannot detach and then reattach instances to an ELB within an Elastic Beanstalk application
D. The instances needed to be reattached before the new application version was deployed

Answer: B

Explanation:
Before the CC2 Instances can start receiving traffic, they will be checked via the health checks of the CLB. Once the health checks are successful, the CC2 Instance
will change its state to InService and then the EC2 Instances can start receiving traffic. 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 139
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 143
You are deciding on a deployment mechanism for your application. Which of the following deployment mechanisms provides the fastest rollback after failure.

A. Rolling-Immutable
B. Canary
C. Rolling-Mutable
D. Blue/Green

Answer: D

Explanation:
In Blue Green Deployments, you will always have the previous version of your application available.
So anytime there is an issue with a new deployment, you can just quickly switch back to the older version of your application.
For more information on Blue Green Deployments, please refer to the below link: https://docs.cloudfoundry.org/devguide/deploy-apps/blue-green.html

NEW QUESTION 146
You currently have an application with an Auto Scalinggroup with an Elastic Load Balancer configured in AWS. After deployment users are complaining of slow response time for your application. Which of the following can be used as a start to diagnose the issue

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

Answer: B

Explanation:
High latency on the ELB side can be caused by several factors, such as:
• Network connectivity
• ELB configuration
• Backend web application server issues
For more information on ELB latency, please refer to the below link:
• https://aws.amazon.com/premiumsupport/knowledge-center/elb-latency-troubleshooting/

NEW QUESTION 149
Which of the below services can be used to deploy application code content stored in Amazon S3 buckets, GitHub repositories, or Bitbucket repositories

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

Answer: B

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

NEW QUESTION 150
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 154
Which of the following is the right sequence of initial steps in the deployment of application revisions using Code Deploy
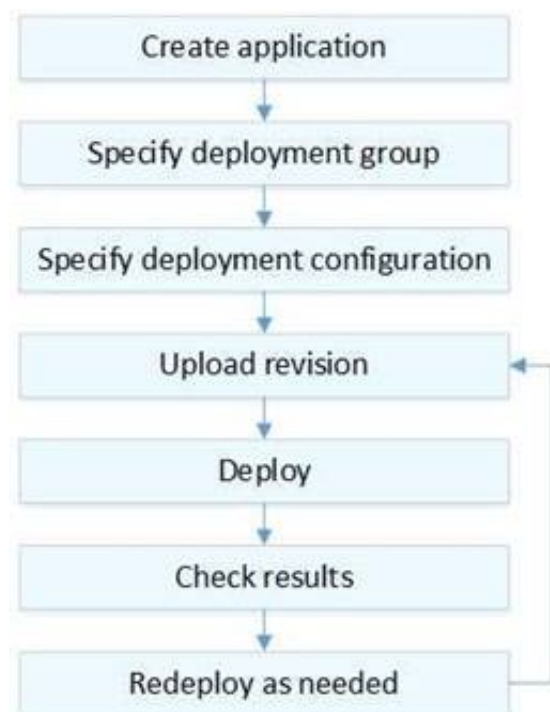1) Specify deployment configuration
2) Upload revision
3) Create application
4) Specify deployment group

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

```
┌─────────────────────────┐
│    Create application    │
└─────────────────────────┘
             │
             ▼
┌─────────────────────────┐
│ Specify deployment group │
└─────────────────────────┘
             │
             ▼
┌──────────────────────────────┐
│ Specify deployment configuration │
└──────────────────────────────┘
             │
             ▼
┌─────────────────────────┐
│     Upload revision     │◄──────┐
└─────────────────────────┘       │
             │                    │
             ▼                    │
┌─────────────────────────┐       │
│         Deploy          │       │
└─────────────────────────┘       │
             │                    │
             ▼                    │
┌─────────────────────────┐       │
│      Check results      │       │
└─────────────────────────┘       │
             │                    │
             ▼                    │
┌─────────────────────────┐       │
│    Redeploy as needed    │──────┘
└─────────────────────────┘
```

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 156
Which of the following services can be used to implement DevOps in your company.

A. AWS Elastic Beanstalk
B. AWSOpswork
C. AWS Cloudformation
D. All of the above

Answer: D

Explanation:
All of the services can be used to implement Devops in your company
1) AWS Elastic Beanstalk, an easy-to-use service for deploying and scaling web applications and services developed with Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker on servers such as Apache, Nginx, Passenger, and I IS.
2) AWS Ops Works, a configuration management service that helps you configure and operate applications of all shapes and sizes using Chef
3) AWS Cloud Formation, which is an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable fashion.
For more information on AWS Devops please refer to the below link:
• http://docs.aws.amazon.com/devops/latest/gsg/welcome.html

NEW QUESTION 160
Which of the following environment types are available in the Elastic Beanstalk environment. Choose 2 answers from the options given below

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

Answer: AC

Explanation:
The AWS Documentation mentions
In Elastic Beanstalk, you can create a load-balancing, autoscaling environment or a single-instance environment. The type of environment that you require depends
on the application that you deploy.
When you go onto the Configuration for your environment, you will be able to see the Environment type from there

NEW QUESTION 164
You have 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

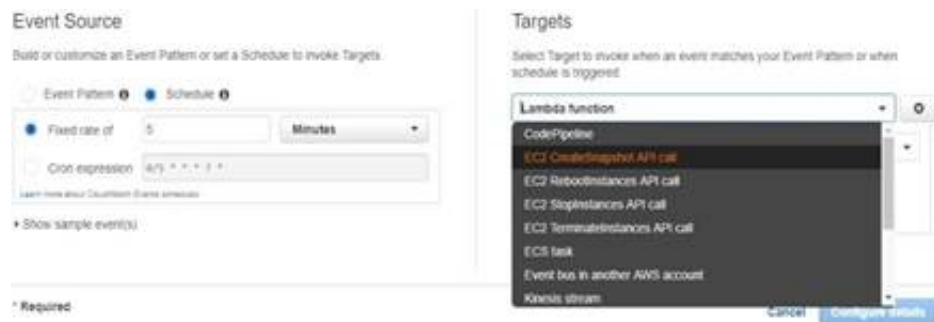**Answer:** D

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/WhatI sCloudWatchCvents.html

NEW QUESTION 169
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 171
You currently have an Autoscalinggroup that has the following settings Min capacity-2
Desired capacity - 2 Maximum capacity - 4
The current number of instances running in the Autoscaling Group is 2. You have been notified that for a duration of an hour, you need to ensure that no new instances are launched by the Autoscaling Group Which of the below 2 actions can be carried out to fulfil this requirement

A. Change the Maximum capacity to 2
B. Change the Desired capacity to 4
C. Suspend the Launch process of the Autoscaling Group
D. Change the Minimum capacity to 2

Answer: AC

Explanation:
You can temporarily suspend the creation of new instances by either reducing the Maximum capacity to 2, so that the current instances running which is 2, matches the maximum limit.
Secondly you can suspend the launch process of the Autoscaling Group The AWS Documentation mentions
Scaling Processes
Amazon CC2 Auto Scaling supports the following scaling processes:
Launch Adds a new CC2 instance to the group, increasing its capacity. Warning
If you suspend Launch, this disrupts other processes. For example, you can't return an instance in a standby state to service if the Launch process is suspended, because the group can't scale. For more information on Autoscaling Suspend and Resume, please visit the below URL:
? http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resume-processes.html

NEW QUESTION 173
Your company is planning to setup a wordpress application. The wordpress application will connect to a MySQL database. Part of the requirement is to ensure that the database environment is fault

tolerant and highly available. Which of the following 2 options individually can help fulfil this requirement.

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

Answer: AD

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

NEW QUESTION 175
You are currently using Elastic Beanstalk to host your production environment. You need to rollout updates to your application hosted on this environment. This is a critical application which is why there is a requirement that the rollback, if required, should be carried out with the least amount of downtime. Which of the following deployment strategies would ideally help achieve this purpose

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

Answer: C

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

NEW QUESTION 176
Your company has an on-premise Active Directory setup in place. The company has extended their footprint on AWS, but still want to have the ability to use their on-premise Active Directory for authentication. Which of the following AWS services can be used to ensure that AWS resources such as AWS Workspaces can continue to use the existing credentials stored in the on-premise Active Directory.

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

Answer: C

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

NEW QUESTION 180
You have a legacy application running that uses an m4.large instance size and cannot scale with Auto Scaling, but only has peak performance 5% of the time. This is a huge waste of resources and money so your Senior Technical Manager has set you the task of trying to reduce costs while still keeping the legacy application running as it should. Which of the following would best accomplish the task your manager has set you? Choose the correct answer from the options below

A. Use a T2burstable performance instance.
B. Use a C4.large instance with enhanced networking.
C. Use two t2.nano instances that have single Root I/O Visualization.
D. Use t2.nano instance and add spot instances when they are required.

Answer: A

Explanation:
The aws documentation clearly indicates using T2 CC2 instance types for those instances which don't use CPU that often.
T2
T2 instances are Burstable Performance Instances that provide a baseline level of CPU performance with the ability to burst above the baseline.

T2 Unlimited instances can sustain high CPU performance for as long as a workload needs it. For most general-purpose workloads, T2 Unlimited instances will provide ample performance without any additional charges. If the instance needs to run at higher CPU utilization for a prolonged period, it can also do so at a flat additional charge of 5 cents per vCPU-hour.

The baseline performance and ability to burst are governed by CPU Credits. T2 instances receive CPU Credits continuously at a set rate depending on the instance size, accumulating CPU Credits when they are idle, and consuming CPU credits when they are active. T2 instances are a good choice for a variety of general-purpose workloads including micro-services, low-latency interactive applications, small and medium databases, virtual desktops, development, build and stage environments, code repositories, and product prototypes. For more information see Burstable Performance Instances.

For more information on F_C2 instance types please see the below link: https://aws.amazon.com/ec2/instance-types/

NEW QUESTION 183

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 188

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

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

Answer: A

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

NEW QUESTION 191

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

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

Answer: BC

Explanation:
Using User data and cloud-init directives you can always ensure you download the latest version of the LAMP stack and give it to the development teams. With AMI's
you will always have the same version and will need to create an AMI everytime the version of the LAMP stack changes.
The AWS Documentation mentions
When you launch an instance in Amazon CC2, you have the option of passing user data to the instance that can be used to perform common automated configuration tasks and even run scripts after the instance starts. You can pass two types of user data to Amazon CC2: shell scripts and cloud-init directives. You can
also pass this data into the launch wizard as plain text, as a file (this is useful for launching instances using the command line tools), or as base64-encoded text (for API calls).
For more information on User data please refer to the below link: http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/user-data.html
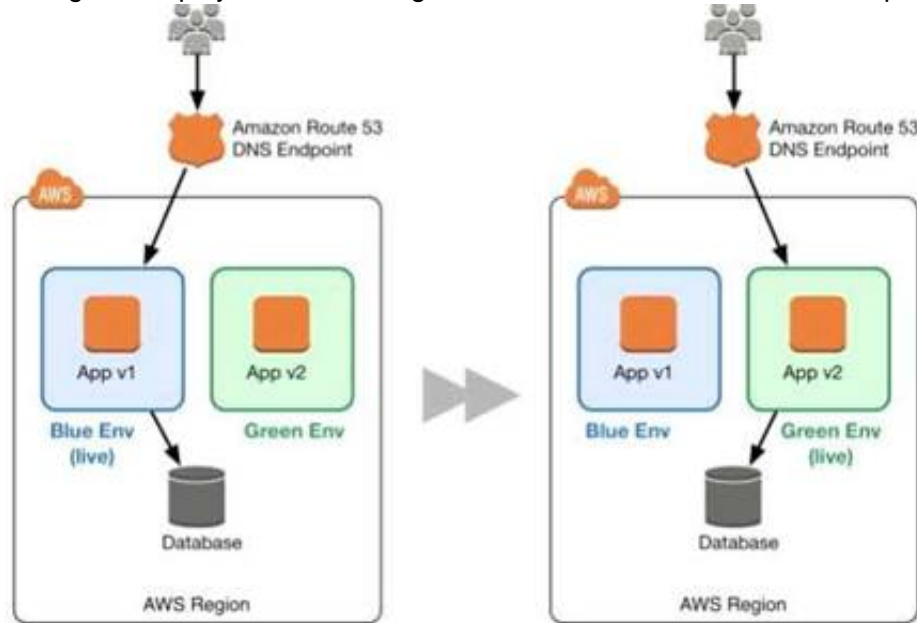
NEW QUESTION 192

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

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

Answer: B

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

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

NEW QUESTION 195
You are in charge of designing Cloudformation templates for your company. One of the key requirements is to ensure that if a Cloudformation stack is deleted, a snapshot of the relational database is created which is part of the stack. How can you achieve this in the best possible way?

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

Answer: C

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

NEW QUESTION 197
Your company is 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 202

You are using lifecycle hooks in your AutoScaling Group. Because there is a lifecycle hook, the instance is put in the Pending:Wait state, which means that it is not available to handle traffic yet. When the instance enters the wait state, other scaling actions are suspended. After some time, the instance state is changed to Pending:Proceed, and finally InService where the instances that are part of the Autoscaling Group can start serving up traffic. But you notice that the bootstrapping process on the instances finish much earlier, long before the state is changed to PendingiProceed.
What can you do to ensure the instances are placed in the right state after the bootstrapping process is complete?

A. Use the complete-lifecycle-action call to complete the lifecycle actio
B. Run this command from another EC2 Instance.
C. Use the complete-lifecycle-action call to complete the lifecycle actio
D. Run this command from the Command line interfac
E. -^C Use the complete-lifecycle-action call to complete the lifecycle actio
F. Run this command from the Simple Notification service.
G. Use the complete-lifecycle-action call to complete the lifecycle actio
H. Run this command from a SQS queue

Answer: B

Explanation:
The AWS Documentation mentions the following
3. If you finish the custom action before the timeout period ends, use the complete-1ifecycle-action command so that the Auto Scalinggroup can continue launching
or terminating the instance. You can specify the lifecycle action token, as shown in the following command:
3. If you finish the custom action before the timeout period ends, use the complete-lifecycle-action command so that Auto Scaling can continue launching or terminating the instance. You can specify the lifecycle action token, as shown in the following command:


```
aws autoscaling complete-lifecycle-action --lifecycle-action-result CONTINUE
```

For more information on lifecycle hooks, please refer to the below URL:
• http://docs.aws.amazon.com/autoscaling/latest/userguide/lifecycle-hooks.htm I

NEW QUESTION 203
There is a requirement for an application hosted on a VPC to access the On-premise LDAP server. The VPC and the On-premise location are connected via an I PSec VPN. Which of the below are the right options for the application to authenticate each user. Choose 2 answers from the options below

A. Develop an identity broker that authenticates against 1AM security Token service to assume a 1AM role in order to get temporary AWS security credentials The application calls the identity broker to get AWS temporary security credentials.
B. The application authenticates against LDAP and retrieves the name of an 1AM role associated with the use
C. The application then calls the 1AM Security Token Service to assume that 1AM rol
D. The application can use the temporary credentials to access any AWS resources.
E. Develop an identity broker that authenticates against LDAP and then calls 1AM Security Token Service to get 1AM federated user credential
F. The application calls the identity broker to get 1AM federated user credentials with access to the appropriate AWS service.
G. The application authenticates against LDAP the application then calls the AWS identity and Access Management (1AM) Security service to log in to 1AM using the LDAP credentials the application can use the 1AM temporary credentials to access the appropriate AWS service.
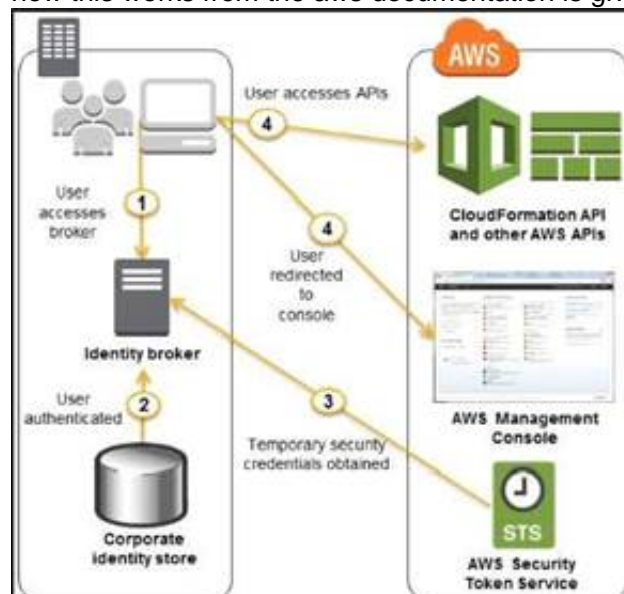
Answer: BC

Explanation:
When you have the need for an in-premise environment to work with a cloud environment, you would normally have 2 artefacts for authentication purposes
• An identity store - So this is the on-premise store such as Active Directory which stores all the information for the user's and the groups they below to.
• An identity broker - This is used as an intermediate agent between the on-premise location and the cloud environment. In Windows you have a system known as Active Directory Federation services to provide this facility.
Hence in the above case, you need to have an identity broker which can work with the identity store and the Security Token service in aws. An example diagram of how this works from the aws documentation is given below.

For more information on federated access, please visit the below link: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_common-scenarios_federated-users.htm I

NEW QUESTION 208
......

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