DOP-C01 Dumps

AWS Certified DevOps Engineer- Professional

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

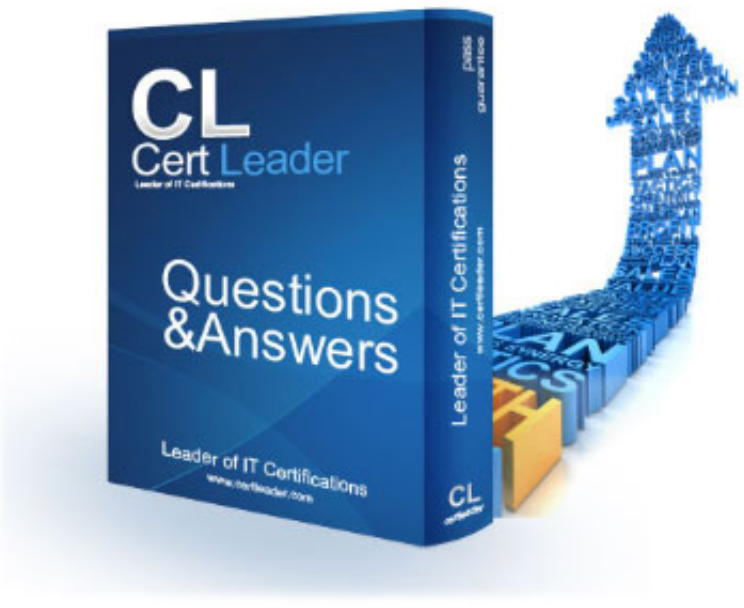

NEW QUESTION 1
Your application stores sensitive information on an EBS volume attached to your EC2 instance. How can you protect your information? Choose two answers from the options given below

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 2
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 3
The project you are working on currently uses a single AWS CloudFormation template to deploy its AWS infrastructure, which supports a multi-tier web application. You have been tasked with organizing the AWS CloudFormation resources so that they can be maintained in the future, and so that different departments such as Networking and Security can review the architecture before it goes to Production. How should you do this in a way that accommodates each department, using their existing workflows?

A. Organize the AWS CloudFormation template so that related resources are next to each other in the template, such as VPC subnets and routing rules for Networkingand security groups and 1AM information for Security.
B. Separate the AWS CloudFormation template into a nested structure that has individual templates for the resources that are to be governed by different departments, and use the outputs from the networking and security stacks for the application template that you contro
C. ^/
D. Organize the AWS CloudFormation template so that related resources are next to each other in the template for each department's use, leverage your existing continuous integration tool to constantly deploy changes from all parties to the Production environment, and then run tests for validation.
E. Use a custom application and the AWS SDK to replicate the resources defined in the current AWS CloudFormation template, and use the existing code review system to allow other departments to approve changes before altering the application for future deployments.

Answer: B

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

NEW QUESTION 4

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 5
As part of your continuous deployment process, your application undergoes an I/O load performance test before it is deployed to production using new AMIs. The application uses one Amazon Elastic Block Store (EBS) PIOPS volume per instance and requires consistent I/O performance. Which of the following must be carried out to ensure that I/O load performance tests yield the correct results in a repeatable manner?

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

Answer: B

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

NEW QUESTION 6
Your development team wants account-level access to production instances in order to do live debugging of a highly secure environment. Which of the following should you do?

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

Answer: C

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

NEW QUESTION 7
When an Auto Scaling group is running in Amazon Elastic Compute Cloud (EC2), your application rapidly scales up and down in response to load within a 10-minute window; however, after the load peaks, you begin to see problems in your configuration management system where previously terminated Amazon EC2 resources are still showing as active. What would be a reliable and efficient way to handle the cleanup of Amazon EC2 resources within your configuration management system? Choose two answers from the options given below

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

Answer: AD

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

NEW QUESTION 8
You have been tasked with deploying a scalable distributed system using AWS OpsWorks. Your distributed system is required to scale on demand. As it is distributed, each node must hold a configuration file that includes the hostnames of the other instances within the layer. How should you configure AWS OpsWorks to manage scaling this application dynamically?

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

Answer: A

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

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

NEW QUESTION 9
You have a web application that's developed in Node.js The code is hosted in Git repository. You want to now deploy this application to AWS. Which of the below 2 options can fulfil this requirement.

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

Answer: CD

Explanation:

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

NEW QUESTION 10
You are using Elastic Beanstalk to manage your application. You have a SQL script that needs to only be executed once per deployment no matter how many EC2 instances you have running. How can you do this?

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

Answer: C

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

NEW QUESTION 10
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.

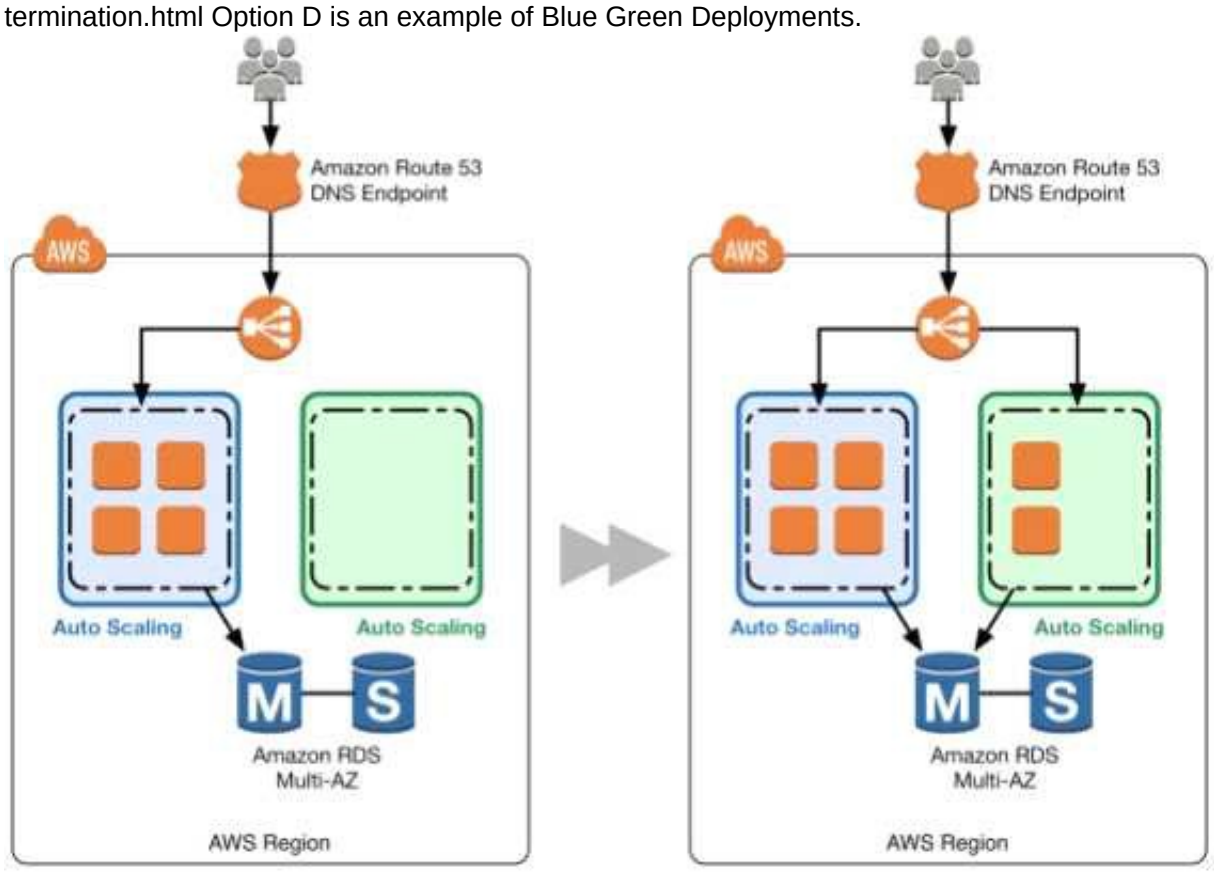

A blue group carries the production load while a green group is staged and deployed with the new code. When if s time to deploy, you simply attach the green group to the existing load balancer to introduce traffic to the new environment. For HTTP/HTTP'S listeners, the load balancer favors the green Auto 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 14
You have a development team that is continuously spending a lot of time rolling back updates for an application. They work on changes, and if the change fails, they spend more than 5-6h in rolling back the update. Which of the below options can help reduce the time for rolling back application versions.

A. Use Elastic Beanstalk and re-deploy using Application Versions
B. Use S3 to store each version and then re-deploy with Elastic Beanstalk
C. Use CloudFormation and update the stack with the previous template
D. Use OpsWorks and re-deploy using rollback feature.

Answer: A

Explanation:
Option B is invalid because Clastic Beanstalk already has the facility to manage various versions and you don't need to use S3 separately for this.
Option C is invalid because in Cloudformation you will need to maintain the versions. Clastic Beanstalk can so that automatically for you.
Option D is good for production scenarios and Clastic Beanstalk is great for development scenarios. AWS beanstalk is the perfect solution for developers to maintain application versions.
With AWS Clastic Beanstalk, you can quickly deploy and manage applications in the AWS Cloud without worrying about the infrastructure that runs those applications. AWS Clastic Beanstalk reduces management complexity without restricting choice or control. You simply upload your application, and AWS Clastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring.
For more information on AWS Beanstalk please refer to the below link: https://aws.amazon.com/documentation/elastic-beanstalk/

NEW QUESTION 19
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 23
You have an application hosted in AWS. You wanted to ensure that when certain thresholds are reached, a Devops Engineer is notified. Choose 3 answers from the options given below

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

Answer: ACD

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

NEW QUESTION 28
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 31
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.html
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 34
You are planning on using the Amazon RDS facility for Fault tolerance for your application. How does Amazon RDSmuIti Availability Zone model work

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

Answer: A

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

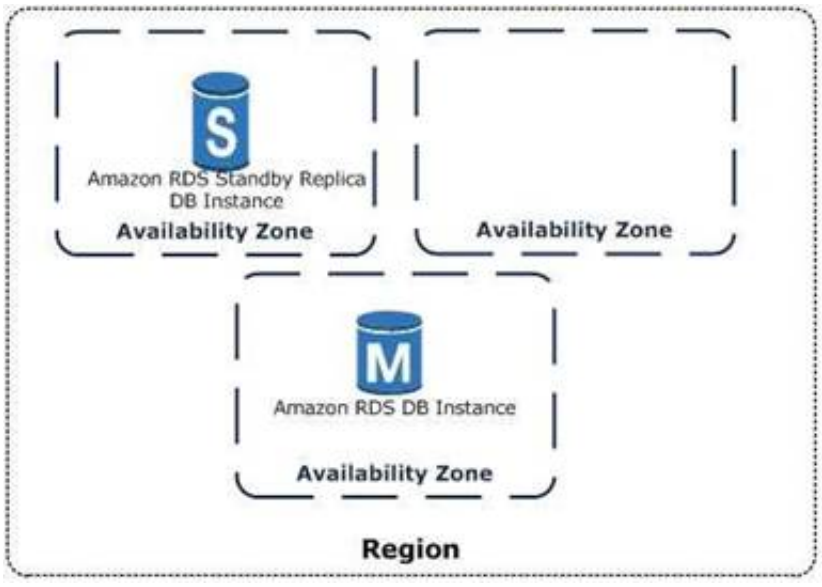

Option B is invalid because the replication is synchronous.
Option C and D are invalid because this is built around AZ and not regions. For more information on Multi-AZ RDS, please visit the below URL:
https://aws.amazon.com/rds/details/multi-az/

NEW QUESTION 38
Your system automatically provisions EIPs to EC2 instances in a VPC on boot. The system provisions the whole VPC and stack at once. You have two of them per VPC. On your new AWS account, your attempt to create a Development environment failed, after successfully creating Staging and Production environments in the same region. What happened?

A. You didn't choose the Development version of the AMI you are using.
B. You didn't set the Development flag to true when deploying EC2 instances.
C. You hit the soft limit of 5 EIPs per region and requested a 6th.
D. You hit the soft limit of 2 VPCs per region and requested a 3rd.

Answer: C

Explanation:
The most likely cause is the fact you have hit the maximum of 5 Elastic IP's per region.
By default, all AWS accounts are limited to 5 Clastic IP addresses per region, because public (IPv4) Internet addresses are a scarce public resource. We strongly encourage you to use an Elastic IP address primarily for the ability to remap the address to another instance in the case of instance failure, and to use DNS hostnames for all other inter-node communication.
Option A is invalid because a AMI does not have a Development version tag. Option B is invalid because there is no flag for an CC2 Instance
Option D is invalid because there is a limit of 5 VPCs per region. For more information on Clastic IP's, please visit the below URL:
• http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/elastic-i p-addresses-eip.html

NEW QUESTION 40
You are designing a service that aggregates clickstream data in batch and delivers reports to subscribers via email only once per week. Data is extremely spikey, geographically distributed, high- scale, and unpredictable. How should you design this system?

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

Answer: B

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

NEW QUESTION 42
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 43
For AWS Auto Scaling, what is the first transition state an existing instance enters after leaving Standby state?

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

Answer: C

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

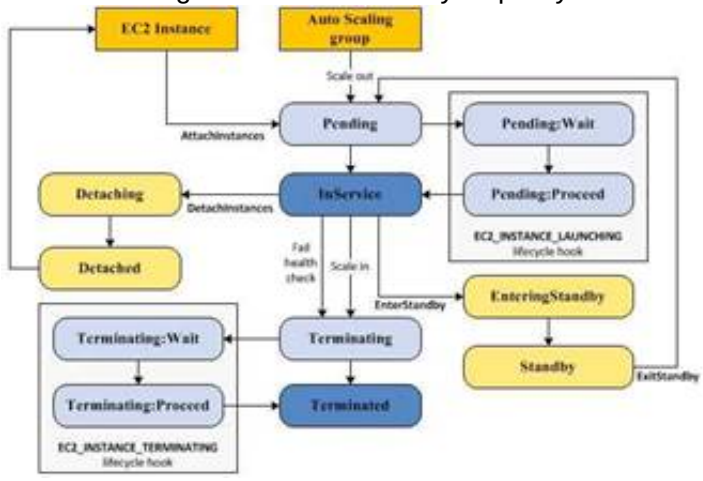

For more information on Autoscaling Lifecycle, please refer to the below link:
http://docs.aws.amazon.com/autoscaling/latest/userguide/AutoScaingGroupLifecycle.html

NEW QUESTION 44
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 46
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 50
What is required to achieve gigabit network throughput on EC2? You already selected cluster- compute, 10GB instances with enhanced networking, and your workload is already network-bound, but you are not seeing 10 gigabit speeds.

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

Answer: D

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

NEW QUESTION 51
You need to deploy a new application version to production. Because the deployment is high-risk, you need to roll the new version out to users over a number of hours, to make sure everything is working correctly. You need to be able to control the proportion of users seeing the new version of the application down to the percentage point. You use ELB and EC2 with Auto Scaling Groups and custom AMIs with your code pre-installed assigned to Launch Configurations. There are no data base- level changes during your deployment. You have been told you cannot spend too much money, so you must not increase the number of EC2 instances much at all during the deployment, but you also need to be able to switch back to the original version of code quickly if something goes wrong. What is the best way to meet these requirements?

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

Answer: A

Explanation:
This is an example of a Blue Green Deployment
You can shift traffic all at once or you can do a weighted distribution. With Amazon Route 53, you can define a percentage of traffic to go to the green environment and gradually update the weights until the green environment carries the full production traffic. A weighted distribution provides the ability to perform canary analysis where a small percentage of production traffic is introduced to a new environment. You can test the new code and monitor for errors, limiting the blast radius if any issues are encountered. It also allows the green environment to scale out to support the full production load if you're using Elastic Load Balancing

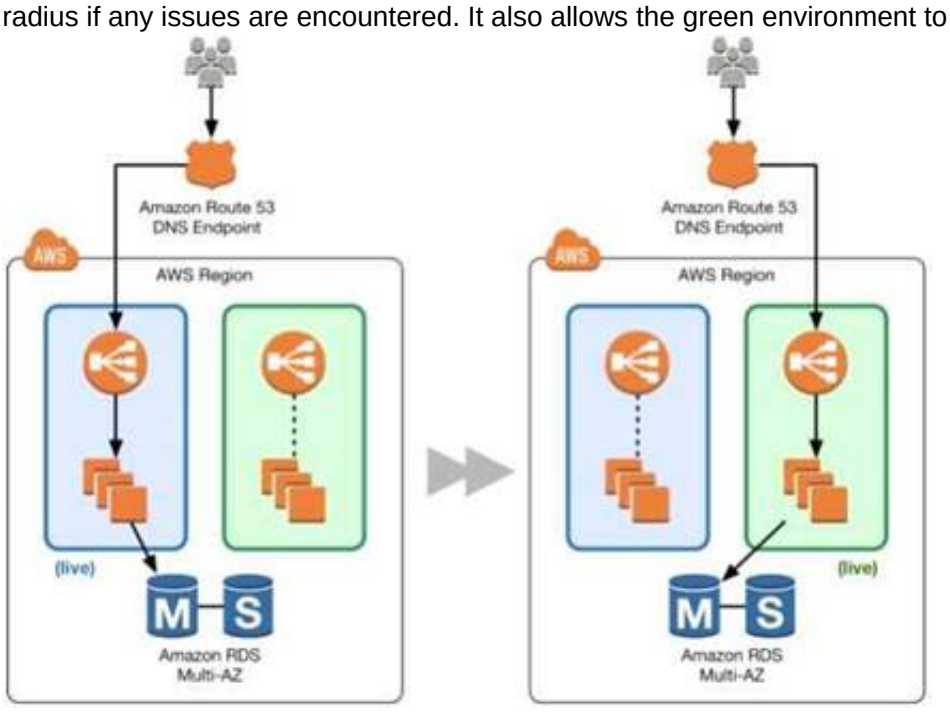

For more information on Blue Green Deployments, please visit the below URL:
• https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 54
You are building a mobile app for consumers to post cat pictures online. You will be storing the images in AWS S3. You want to run the system very cheaply and

simply. Which one of these options allows you to build a photo sharing application with the right authentication/authorization implementation.

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

Answer: A

Explanation:
Amazon Cognito lets you easily add user sign-up and sign-in and manage permissions for your mobile and web apps. You can create your own user directory within Amazon Cognito. You can also choose to authenticate users through social identity providers such as Facebook, Twitter, or Amazon; with SAML identity solutions; or by using your own identity system. In addition, Amazon Cognito enables you to save data locally on users' devices, allowing your applications to work even when the devices are offline. You can then synchronize data across users' devices so that their app experience remains consistent regardless of the device they use.
For more information on AWS Cognito, please visit the below URL:
• http://docs.aws.amazon.com/cognito/latest/developerguide/what-is-amazon-cognito.html

NEW QUESTION 57
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 58
Your CTO thinks your AWS account was hacked. What is the only way to know for certain if there was unauthorized access and what they did, assuming your hackers are very sophisticated AWS engineers and doing everything they can to cover their tracks?

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

Answer: A

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

NEW QUESTION 61
You have an application hosted in AWS, which sits on EC2 Instances behind an Elastic Load Balancer. You have added a new feature to your application and are now receving complaints from users that the site has a slow response. Which of the below actions can you carry out to help you pinpoint the issue

A. Use Cloudtrail to log all the API calls, and then traverse the log files to locate the issue
B. Use Cloudwatch, monitor the CPU utilization to see the times when the CPU peaked
C. Reviewthe Elastic Load Balancer logs
D. Create some custom Cloudwatch metrics which are pertinent to the key features of your application

Answer: D

Explanation:
Since the issue is occuring after the new feature has been added, it could be relevant to the new feature.
Enabling Cloudtrail will just monitor all the API calls of all services and will not benefit the cause.
The monitoring of CPU utilization will just reverify that there is an issue but will not help pinpoint the issue.
The Elastic Load Balancer logs will also just reverify that there is an issue but will not help pinpoint the issue.
For more information on custom Cloudwatch metrics, please refer to the below link:
http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.html

NEW QUESTION 66
When creating an Elastic Beanstalk environment using the Wizard, what are the 3 configuration options presented to you

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

Answer: ABD

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

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

NEW QUESTION 71
Which of the following features of the Elastic Beanstalk service will allow you to perform a Blue Green Deployment

A. Rebuild Environment
B. Swap Environment
C. Swap URL's
D. Environment Configuration

Answer: C

Explanation:
With the Swap url feature, you can keep a version of your environment ready. And when you are ready to cut over, you can just use the swap url feature to switch over
to your new environment
For more information on swap url feature, please refer to the below link:
• http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAMCSwap.html

NEW QUESTION 76
You have an Autoscaling Group configured to launch EC2 Instances for your application. But you notice that the Autoscaling Group is not launching instances in the right proportion. In fact instances are being launched too fast. What can you do to mitigate this issue? Choose 2 answers from the options given below

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

Answer: AB

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

NEW QUESTION 78
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 82
There is a requirement for a vendor to have access to an S3 bucket in your account. The vendor already has an AWS account. How can you provide access to the vendor on this bucket.

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

Answer: C

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

NEW QUESTION 84
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 86
You need to deploy a multi-container Docker environment on to Elastic beanstalk. Which of the following files can be used to deploy a set of Docker containers to Elastic beanstalk

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

Answer: C

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

NEW QUESTION 91
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 95
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 98
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 103
You are incharge of creating a Cloudformation template that will be used to spin our resources on demand for your Devops team. The requirement is that this cloudformation template should be able to spin up resources in different regions. Which of the following aspects of Cloudformation templates can help you design the template to spin up resources based on the region.

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

Answer: A

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

NEW QUESTION 106

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 107
Your company is planning to develop an application in which the front end is in .Net and the backend is in DynamoDB. There is an expectation of a high load on the application. How could you ensure the scalability of the application to reduce the load on the DynamoDB database? Choose an answer from the options below.

A. Add more DynamoDB databases to handle the load.
B. Increase write capacity of Dynamo DB to meet the peak loads
C. Use SQS to assist and let the application pull messages and then perform the relevant operation in DynamoDB.
D. Launch DynamoDB in Multi-AZ configuration with a global index to balance writes

Answer: C

Explanation:
When the idea comes for scalability then SQS is the best option. Normally DynamoDB is scalable, but since one is looking for a cost effective solution, the messaging in SQS can assist in managing the situation mentioned in the question.
Amazon Simple Queue Service (SQS) is a fully-managed message queuing service for reliably communicating among distributed software components and microservices - at any scale. Building applications from individual components that each perform a discrete function improves scalability and reliability, and is best practice design for modern applications. SQS makes it simple and cost- effective to decouple and coordinate the components of a cloud application. Using SQS, you can send, store, and receive messages between software components at any volume, without losing messages or requiring other services to be always available
For more information on SQS, please refer to the below URL:
• https://aws.amazon.com/sqs/

NEW QUESTION 111
Which of the following features of the Autoscaling Group ensures that additional instances are neither launched or terminated before the previous scaling activity takes effect

A. Termination policy
B. Cool down period
C. Ramp up period
D. Creation policy

Answer: B

Explanation:
The AWS documentation mentions
The Auto Scaling cooldown period is a configurable setting for your Auto Scaling group that helps to ensure that Auto Scaling doesn't launch or terminate additional
instances before the previous scaling activity takes effect. After the Auto Scaling group dynamically scales using a simple scaling policy. Auto Scaling waits for the cooldown period to complete before resuming scaling activities. When you manually scale your Auto Scaling group, the default is not to wait for the cooldown period,
but you can override the default and honor the cooldown period. If an instance becomes unhealthy.
Auto Scaling does not wait for the cooldown period to complete before replacing the unhealthy instance
For more information on the Cool down period, please refer to the below URL:
• http://docs.ws.amazon.com/autoscaling/latest/userguide/Cooldown.htmI

NEW QUESTION 115
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 119
You are using Autoscaling for managing the instances in your AWS environment. You need to deploy a new version of your application. You'd prefer to use all new instances if possible, but you cannot have any downtime. You also don't want to swap any environment urls. Which of the following deployment methods would you implement

A. Using "All at once" deployment method.
B. Using "Blue Green" deployment method.
C. Using"RollingUpdates"deploymentmethod.
D. Using "Blue Green" with "All at once" deployment method.

Answer: C

Explanation:
In Rolling deployment, you can mention a new set of servers which can replace the existing set of servers. This replacement will happen in a phased out manner. Since there is a requirement to not swap URL's, you must not use Blue Green deployments.
For more information on the differences between Rolling Updates and Blue Green deployments, please refer to the below URL:
• https://cloudnative.io/docs/blue-green-deployment/

NEW QUESTION 124
Your IT company is currently hosting a production environment in Elastic beanstalk. You understand that the Elastic beanstalk service provides a facility known as Managed updates which are minor and patch version updates which are periodically required for your system. Your IT supervisor is worried about the impact that these updates would have on the system. What can you tell about the Elastic beanstalk service with regards to managed updates

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

Answer: D

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

NEW QUESTION 129
A user is trying to save some cost on the AWS services. Which of the below mentioned options will not help him save cost?

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

Answer: B

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

NEW QUESTION 132
Which of the following are components of the AWS Data Pipeline service. Choose 2 answers from the options given below

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

Answer: AB

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

NEW QUESTION 135
Which of the following are Lifecycle events available in Opswork? Choose 3 answers from the options below

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

Answer: ACD

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

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

NEW QUESTION 137
You are writing an AWS Cloud Formation template and you want to assign values to properties that will not be available until runtime. You know that you can use intrinsic functions to do this but are unsure as to which part of the template they can be used in. Which of the following is correct in describing how you can currently use intrinsic functions in an AWS CloudFormation template?

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

Answer: B

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

NEW QUESTION 140
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. DeployElasticCache in-memory cache running in each availability zone
B. Implementshardingto distribute load to multiple RDS MySQL instances
C. Increasethe RDS MySQL Instance size and Implement provisioned IOPS
D. Addan RDS MySQL read replica in each availability zone

Answer: AD

Explanation:

Implement Read Replicas and Elastic 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 145
Which of the below 3 things can you achieve with the Cloudwatch logs service? Choose 3 options.

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

Answer: BCD

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

NEW QUESTION 150
Your company is using an Autoscaling Group to scale out and scale in instances. There is an expectation of a peak in traffic every Monday at 8am. The traffic is then expected to come down before the weekend on Friday 5pm. How should you configure Autoscaling in this?

A. Createdynamic scaling policies to scale up on Monday and scale down on Friday
B. Create a scheduled policy to scale up on Fridayand scale down on Monday
C. CreateascheduledpolicytoscaleuponMondayandscaledownonFriday
D. Manuallyadd instances to the Autoscaling Group on Monday and remove them on Friday

Answer: C

Explanation:
The AWS Documentation mentions the following for Scheduled scaling
Scaling based on a schedule allows you to scale your application in response to predictable load changes. For example, every week the traffic to your web application starts to increase on Wednesday, remains high on Thursday, and starts to decrease on Friday. You can plan your scaling activities based on the predictable traffic patterns of your web application.
For more information on scheduled scaling for Autoscaling, please visit the below URL
• http://docs.aws.amazon.com/autoscaling/latest/userguide/sched ule_time.htm I

NEW QUESTION 152
Which of the following is not a supported platform on Elastic Beanstalk?

A. PackerBuilder
B. Go
C. Nodejs
D. JavaSE
E. Kubernetes

Answer: E

Explanation:
Answer-C
Below is the list of supported platforms
*Packer Builder
*Single Container Docker
*Multicontainer Docker
*Preconfigured Docker
*Go
*Java SE
*Java with Tomcat
*NET on Windows Server with I IS
*Nodejs
*PHP
*Python
*Ruby
For more information on the supported platforms please refer to the below link
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/concepts.platforms. Html

NEW QUESTION 156
Your CTO is very worried about the security of your AWS account. How best can you prevent hackers from completely hijacking your account?

A. Useshort but complex password on the root account and any administrators.
B. UseAWS 1AM Geo-Lock and disallow anyone from logging in except for in your city.
C. UseMFA on all users and accounts, especially on the root account.
D. Don'twrite down or remember the root account password after creating the AWSaccount.

Answer: C

Explanation:
The AWS documentation mentions the following on MFA
AWS Multi-Factor Authentication (MFA) is a simple best practice that adds an extra layer of protection on top of your user name and password. With MFA enabled, when a user signs in to an AWS website, they will be prompted for their user name and password (the first factor—what they know), as well as for an authentication code from their AWS MFA device (the second factor—what they have). Taken together, these multiple factors provide increased security for your AWS account settings and resources.
For more information on MFA please visit the below link https://aws.ama zon.com/ia m/detai Is/mfa/

NEW QUESTION 157
How can you resolve a dependency Error when using Cloudformation?

A. Usethe mappings attribute
B. Usethe parameter attribute
C. UsetheDependsOn attribute
D. Use the Error attribute

Answer: C

Explanation:
The AWS troubleshooting guide for Cloudformation states the following
To resolve a dependency error, add a DependsOn attribute to resources that depend on other resources in your template. In some cases, you must explicitly declare dependencies so that AWS Cloud Formation can create or delete resources in the correct order. For example, if you create an Elastic IP and a VPC with an Internet gateway in the same stack, the Elastic IP must depend on the Internet gateway attachment.
For more information on Cloudformation troubleshooting, please refer to the below url
http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/troubleshooting.html

NEW QUESTION 160
You are in charge of designing a number of Cloudformation templates for your organization. You need to ensure that no one can accidentally update the production based resources on the stack during a stack update. How can this be achieved in the most efficient way?

A. Createtags for the resources and then create 1AM policies to protect the resources.
B. Usea Stack based policy to protect the production based resources.
C. UseS3 bucket policies to protect the resources.
D. UseMFA to protect the resources

Answer: B

Explanation:
The AWS Documentation mentions
When you create a stack, all update actions are allowed on all resources. By default, anyone with stack update permissions can update all of the resources in the stack. During an update, some resources might require an interruption or be completely replaced, resulting in new physical IDs or completely new storage. You can prevent stack resources from being unintentionally updated or deleted during a stack update by using a stack policy. A stack policy is a JSON document that defines the update action1.-; that car1 be performed on designated resources.
For more information on protecting stack resources, please visit the below url http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/protect-stack-resources.html

NEW QUESTION 161
You are Devops Engineer for a large organization. The company wants to start using Cloudformation templates to start building their resources in AWS. You are getting requirements for the templates from various departments, such as the networking, security, application etc. What is the best way to architect these Cloudformation templates.

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

Answer: B

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

NEW QUESTION 162
You are in charge of designing a number of Cloudformation templates for your organization. You are required to make changes to stack resources every now and then based on the requirement. How can you check the impact of the change to resources in a cloudformation stack before deploying changes to the stack?

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

Answer: B

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

NEW QUESTION 166
An application is currently writing a large number of records to a DynamoDB table in one region. There is a requirement for a secondary application tojust take in the changes to the DynamoDB table every 2 hours and process the updates accordingly. Which of the following is an ideal way to ensure the secondary application can get the relevant changes from the DynamoDB table.

A. Inserta timestamp for each record and then scan the entire table for the timestamp asper the last 2 hours.
B. Createanother DynamoDB table with the records modified in the last 2 hours.
C. UseDynamoDB streams to monitor the changes in the DynamoDB table.
D. Transferthe records to S3 which were modified in the last 2 hours

Answer: C

Explanation:
The AWS Documentation mentions the following
A DynamoDB stream is an ordered flow of information about changes to items in an Amazon DynamoDB table. When you enable a stream on a table, DynamoDB captures information about every modification to data items in the table.
Whenever an application creates, updates, or deletes items in the table, DynamoDB Streams writes a stream record with the primary key attribute(s) of the items that were modified. Astream record contains information about a data modification to a single item in a DynamoDB table. You can configure the stream so that the stream records capture additional information, such as the "before" and "after" images of modified items.
For more information on DynamoDB streams, please visit the below URL: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.html

NEW QUESTION 168
Your application is having a very high traffic, so you have enabled autoscaling in multi availability zone to suffice the needs of your application but you observe that one of the availability zone is not receiving any traffic. What can be wrong here?

A. Autoscalingonly works for single availability zone
B. Autoscalingcan be enabled for multi AZ only in north Virginia region
C. Availabilityzone is not added to Elastic load balancer
D. Instancesneed to manually added to availability zone

Answer: C

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

NEW QUESTION 173
You have just developed a new mobile application that handles analytics workloads on large scale datasets that are stored on Amazon Redshift. Consequently, the application needs to access Amazon Redshift tables. Which of the below methods would be the best, both practically and security-wise, to access the tables? Choose the correct answer from the options below

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

Answer: D

Explanation:
For access to any AWS service, the ideal approach for any application is to use Roles. This is the first preference. Hence option A and C are wrong.
For more information on 1AM policies please refer to the below link: http://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html
Next for any web application, you need to use web identity federation. Hence option D is the right option. This along with the usage of roles is highly stressed in the AWS documentation.
"When you write such an app, you'll make requests to AWS services that must be signed with an AWS access key. However, we strongly recommend that you do not embed or distribute long-term AWS credentials with apps that a user downloads to a device, even in an encrypted store. Instead, build your app so that it requests temporary AWS security credentials dynamically when needed using web identity federation. The supplied temporary credentials map to an AWS role that has only

the permissions needed to perform the tasks required by the mobile app".
For more information on web identity federation please refer to the below link: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html

NEW QUESTION 177
Your team is responsible for an AWS Elastic Beanstalk application. The business requires that you move to a continuous deployment model, releasing updates to the application multiple times per day with zero downtime. What should you do to enable this and still be able to roll back almost immediately in an emergency to the previous version?

A. Enablerolling updates in the Elastic Beanstalk environment, setting an appropriatepause time for application startup.
B. Createa second Elastic Beanstalk environment running the new application version, andswap theenvironment CNAMEs.
C. Developthe application to poll for a new application version in your code repository;download and install to each running Elastic Beanstalk instance.
D. Createa second Elastic Beanstalk environment with the new application version, andconfigure the old environment to redirect clients, using the HTTP 301 responsecode, to the new environment

Answer: B

Explanation:
The AWS Documentation mentions the below
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 For more information on Elastic beanstalk swap URL please see the below link:
• http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAM CSwap.html

NEW QUESTION 181
Your social media marketing application has a component written in Ruby running on AWS Elastic Beanstalk. This application component posts messages to social media sites in support of various marketing campaigns. Your management now requires you to record replies to these social media messages to analyze the effectiveness of the marketing campaign in comparison to past and future efforts. You've already developed a new application component to interface with the social media site APIs in order to read the replies. Which process should you use to record the social media replies in a durable data store that can be accessed at any time for analytics of historical data?

A. Deploythe new application component in an Auto Scaling group of Amazon EC2 instances,read the data from the social media sites, store it with Amazon Elastic BlockStore, and use AWS Data Pipeline to publish it to Amazon Kinesis for analytics.
B. Deploythe new application component as an Elastic Beanstalk application, read thedata from the social media sites, store it in DynamoDB, and use Apache Hivewith Amazon Elastic MapReduce for analytics.
C. Deploythe new application component in an Auto Scaling group of Amazon EC2 instances,read the data from the social media sites, store it in Amazon Glacier, and useAWS Data Pipeline to publish it to Amazon RedShift for analytics.
D. Deploythe new application component as an Amazon Elastic Beanstalk application, readthe data from the social media site, store it with Amazon Elastic Block store,and use Amazon Kinesis to stream the data to Amazon Cloud Watch for analytics

Answer: B

Explanation:
The AWS Documentation mentions the below
Amazon DynamoDB is a fast and flexible NoSQL database sen/ice for all applications that need consistent, single-digit millisecond latency at any scale. It is a fully managed cloud database and supports both document and key-value store models. Its flexible data model, reliable performance, and automatic scaling of throughput capacity, makes it a great fit for mobile, web, gaming, ad tech, loT, and many other applications.
For more information on AWS DynamoDB please see the below link:
• https://aws.amazon.com/dynamodb/

NEW QUESTION 185
You are managing an application that contains Go as the front end, MongoDB for document management and is hosted on a relevant Web server. You pre-bake AMI'S with the latest version of the Web server, then user the User Data section to setup the application. You now have a change to the underlying Operating system version and need to deploy that accordingly. How can this be done
in the easiest way possible.

A. Createa new EBS Volume with the relevant OS patches and attach it to the EC2lnstance.
B. Createa Cloudformation stack with the new AMI and then deploy the applicationaccordingly.
C. Createa new pre-baked AM I with the new OS and use the User Data seciton to deploy theapplication.
D. Createan Opswork stack with the new AMI and then deploy the application accordingly.

Answer: C

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

NEW QUESTION 186
You are a Devops Engineer for your company. You are in charge of an application that uses EC2, ELB and Autoscaling. You have been requested to get the ELB access logs. When you try to access the logs, you can see that nothing has been recorded in S3. Why is this the case?

A. Youdon't have the necessary access to the logs generated by ELB.
B. Bydefault ELB access logs are disabled.
C. TheAutoscaling service is not sending the required logs to ELB
D. TheEC2 Instances are not sending the required logs to ELB

Answer: B

Explanation:
The AWS Documentation mentions
Access logging is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logging for your load balancer. Clastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify. You can disable access logging at any time.
For more information on L~LB access logs please see the below link:
• http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/access-log-collection.html

NEW QUESTION 191
You have configured the following AWS services in your organization - Auto Scalinggroup, Elastic Load Balancer, and EC2 instances. You have been requested to terminate an instance from the Autoscaling Group when the CPU utilization is less than 30%. How can you achieve this.

A. Createa Cloud Watch alarm to send a notification to SQ
B. SQS can then remove oneinstance from the Autoscaling Group.
C. Createa CloudWatch alarm to send a notification to the Auto Scalinggroup when theaggregated CPU utilization is less than 30% and configure the Auto Scalingpolicy to remove one instance.
D. Createa CloudWatch alarm to send a notification to the EL
E. The ELB can then removeone instance from the Autoscaling Group.
F. Createa CloudWatch alarm to send a notification to the admin tea
G. The admin team canthen manually terminate an instance from the Autnsraline Groun.

Answer: B

Explanation:
The AWS Documentation mentions the following
You should have two policies, one for scaling in (terminating instances) and one for scaling out (launching instances), for each event to monitor. For example, if you want to scale out when the network bandwidth reaches a certain level, create a policy specifying that Auto Scaling should start a certain number of instances to help with your traffic. But you may also want an accompanying policy to scale in by a certain number when the network bandwidth level goes back down
For more information on the scaling plans, please see the below link: http://docs.aws.amazon.com/autoscaling/latest/userguide/sca I ing_plan.html

NEW QUESTION 193
Which of the following is a reliable and durable logging solution to track changes made to your AWS resources?

A. Createa new CloudTrail trail with one new S3 bucket to store the logs and with theglobal services option selecte
B. Use 1AM roles S3 bucket policies and MultiFactor Authentication (MFA) Delete on the S3 bucket that stores your log
C. V
D. Createa new CloudTrail with one new S3 bucket to store the log
E. Configure SNS tosend log file delivery notifications to your management syste
F. Use 1AM rolesand S3 bucket policies on the S3 bucket that stores your logs.
G. Createa new CloudTrail trail with an existing S3 bucket to store the logs and withthe global services option selecte
H. Use S3 ACLs and Multi FactorAuthentication (M FA) Delete on the S3 bucket that stores your logs.
I. Createthree new CloudTrail trails with three new S3 buckets to store the logs one forthe AWS Management console, one for AWS SDKs and one for command line tools.Use 1AM roles and S3 bucket policies on the S3 buckets that store your logs.

Answer: A

Explanation:
AWS Identity and Access Management (1AM) is integrated with AWS CloudTrail, a sen/ice that logs AWS events made by or on behalf of your AWS account. CloudTrail logs authenticated AWS API calls and also AWS sign-in events, and collects this event information in files that are delivered to Amazon S3 buckets.
You need to ensure that all services are included. Hence option B is partially correct.
Option B and D is wrong because it just adds an overhead for having 3 S3 buckets and SNS notifications.
For more information on Cloudtrail, please visit the below URL:
• http://docs.aws.a mazon.com/IAM/latest/UserGuide/cloudtrail-integration.htm I

NEW QUESTION 198
Which of the following is a container for metrics in Cloudwatch?

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

Answer: B

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

NEW QUESTION 203
You work as a Devops Engineer for your company. There are currently a number of environments hosted via Elastic beanstalk. There is a requirement to ensure to ensure that the rollback time for a new version application deployment is kept to a minimal. Which elastic beanstalk deployment method would fulfil this requirement ?

A. Rollingwith additional batch
B. AllatOnce

C. Blue/Green
D. Rolling

Answer: C

Explanation:
The below table from the AWS documentation shows that the least amount of time is spent in rollbacks when it comes to Blue Green deployments. This is because the only thing that needs to be done is for URL's to be swapped.

Deployment Methods

Method	Impact of Failed Deployment	Deploy Time	Zero Downtime	No DNS Change	Rollback Process	Code Deployed To
All at once	Downtime	◷	×	✓	Manual Redeploy	Existing instances
Rolling	Single batch out of service; any successful batches prior to failure running new application version	◷ ◷ †	✓	✓	Manual Redeploy	Existing instances
Rolling with additional batch	Minimal if first batch fails, otherwise, similar to **Rolling**	◷ ◷ ◷ †	✓	✓	Manual Redeploy	New and existing instances
Immutable	Minimal	◷ ◷ ◷ ◷	✓	✓	Terminate New Instances	New instances
Blue/green	Minimal	◷ ◷ ◷ ◷	✓	×	Swap URL	New instances

For more information on Elastic beanstalk deployment strategies, please visit the below URL: http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.deploy-existing-version, html

NEW QUESTION 207
You have an Opswork stack defined with Linux instances. You have executed a recipe, but the execution has failed. What is one of the ways that you can use to diagnose what was the reason why the recipe did not execute correctly.

A. UseAWS Cloudtrail and check the Opswork logs to diagnose the error
B. UseAWS Config and check the Opswork logs to diagnose the error
C. Logintotheinstanceandcheckiftherecipewasproperlyconfigured.
D. Deregisterthe instance and check the EC2 Logs

Answer: C

Explanation:
The AWS Documentation mentions the following
If a recipe fails, the instance will end up in the setup_failed state instead of online. Even though the instance is not online as far as AWS Ops Works Stacks is concerned, the CC2 instance is running and it's often useful to log in to troubleshoot the issue. For example, you can check whether an application or custom cookbook is correctly installed. The AWS Ops Works Stacks built-in support for SSH and RDP login is
available only for instances in the online state.
For more information on Opswork troubleshooting, please visit the below URL: http://docs.aws.amazon.com/opsworks/latest/userguide/troubleshoot-debug-login.htmI

NEW QUESTION 212
A company is building a two-tier web application to serve dynamic transaction-based content. The data tier is leveraging an Online Transactional Processing (OLTP) database. What services should you leverage to enable an elastic and scalable web tier?

A. ElasticLoad Balancing, Amazon EC2, and Auto Scaling
B. ElasticLoad Balancing, Amazon RDS with Multi-AZ, and Amazon S3
C. AmazonRDS with Multi-AZ andAuto Scaling
D. AmazonEC2, Amazon Dynamo DB, and Amazon S3

Answer: A

Explanation:
The question mentioned a scalable web tier and not a database tier. So Option C, D and B are already automated eliminated, since we do not need a database option. The below example shows an Elastic Load balancer connected to 2 CC2 instances connected via Auto Scaling. This is an example of an elastic and scalable web tier. By scalable we mean that the Auto scaling process will increase or decrease the number of CC2 instances as required.

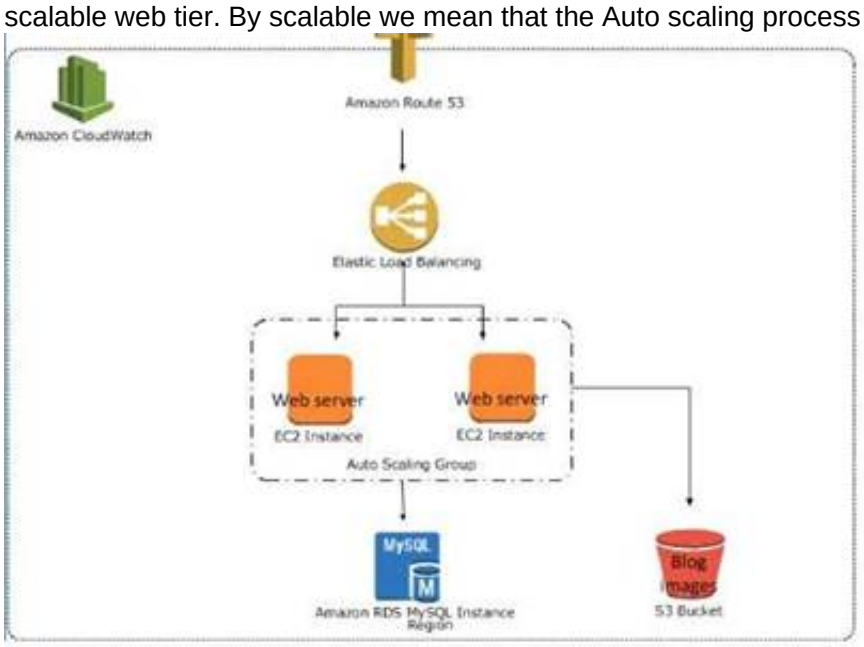

For more information on best practices for AWS Cloud applications, please visit the below URL:
https://dO.awsstatic.com/whitepapers/AWS_Cloud_Best_Practices.pdf

NEW QUESTION 216
You have an ELB on AWS which has a set of web servers behind them. There is a requirement that the SSL key used to encrypt data is always kept secure. Secondly the logs of ELB should only be decrypted by a subset of users. Which of these architectures meets all of the requirements?

A. UseElastic Load Balancing to distribute traffic to a set of web server
B. Toprotect the SSL private key.upload the key to the load balancer and configure the load balancer to offloadthe SSL traffi
C. Write yourweb server logs to an ephemeral volume that has been encrypted using a randomlygenerated AES key.
D. UseElastic Load Balancing to distribute traffic to a set of web server
E. Use TCPIoad balancing on theload balancer and configure your web servers to retrieve the private key from aprivate Amazon S3bucket on boo
F. Write your web server logs to a private Amazon S3 bucket usingAmazon S3 server- sideencryption.
G. UseElastic Load Balancing to distribute traffic to a set of web servers, configurethe load balancer toperform TCP load balancing, use an AWS CloudHSM to perform the SSLtransactions, and write yourweb server logs to a private Amazon S3 bucket using Amazon S3 server-sideencryption.
H. UseElastic Load Balancing to distribute traffic to a set of web server
I. Configurethe load balancer toperform TCP load balancing, use an AWS CloudHSM to perform the SSLtransactions, and write yourweb server logs to an ephemeral volume that has been encrypted using a randomlygenerated AES key.

Answer: C

Explanation:
The AWS CIoudHSM service helps you meet corporate, contractual and regulatory compliance requirements for data security by using dedicated Hardware Security
Module (HSM) appliances within the AWS cloud. With CIoudHSM, you control the encryption keys and cryptographic operations performed by the HSM.
Option D is wrong with the CIoudHSM option because of the ephemeral volume which this is temporary storage
For more information on cloudhsm, please refer to the link:
• https://aws.amazon.com/cloudhsm/

NEW QUESTION 221
A company wants to create standard templates for deployment of their Infrastructure. Which AWS service can be used in this regard? Please choose one option.

A. AmazonSimple Workflow Service
B. AWSEIastic Beanstalk
C. AWSCIoudFormation
D. AWSOpsWorks

Answer: C

Explanation:
AWS Cloud Formation gives developers and systems administrators an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable fashion.
You can use AWS Cloud Formation's sample templates or create your own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run your application. You don't need to figure out the order for provisioning AWS services or the subtleties of making those dependencies work. Cloud Formation takes care of this for you. After the AWS resources are deployed, you can modify and update them in a controlled and predictable way, in effect applying version control to your AWS infrastructure the same way you do with your software. You can also visualize your templates as diagrams and edit them using a drag-and-drop interface with the AWS CloudFormation Designer.
For more information on Cloudformation, please visit the link:
• https://aws.amazon.com/cloudformation/

NEW QUESTION 222
One of your engineers has written a web application in the Go Programming language and has asked your DevOps team to deploy it to AWS. The application code is hosted on a Git repository.
What are your options? (Select Two)

A. Create a new AWS Elastic Beanstalk application and configure a Go environment to host your application, Using Git check out the latest version of the code, once the local repository for Elastic Beanstalk is configured use "eb create" command to create an environment and then use "eb deploy" command to deploy the application.
B. Writea Dockerf ile that installs the Go base image and uses Git to fetch yourapplicatio
C. Create a new AWS OpsWorks stack that contains a Docker layer thatuses the Dockerrun.aws.json file to deploy your container and then use theDockerfile to automate the deployment.
D. Writea Dockerfile that installs the Go base image and fetches your application usingGit, Create a new AWS Elastic Beanstalk application and use this Dockerfile toautomate the deployment.
E. Writea Dockerfile that installs the Go base image and fetches your application usingGit, Create anAWS CloudFormation template that creates and associates an AWS::EC2::lnstanceresource type with an AWS::EC2::Container resource type.

Answer: AC

Explanation:
Opsworks works with Chef recipes and not with Docker containers so Option B and C are invalid. There is no AWS::CC2::Container resource for Cloudformation so Option D is invalid.
Below is the documentation on Clastic beanstalk and Docker
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 Clastic beanstalk and Docker, please visit the link: http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker.html
https://docs.aws.a mazon.com/elasticbeanstalk/latest/dg/eb-cl i3-getting-started.html https://docs^ws.amazon.com/elasticbeanstalk/latest/dg/eb3-cli-githtml

NEW QUESTION 225
The development team has developed a new feature that uses an AWS service and wants to test it from inside a staging VPC. How should you test this feature with the fastest turnaround time?

A. Launchan Amazon Elastic Compute Cloud (EC2) instance in the staging VPC in responseto a development request, and use configuration management to set

up theapplicatio
B. Run any testing harnesses to verify application functionality andthen use Amazon Simple Notification Service (SNS) to notify the developmentteam of the results.
C. Usean Amazon EC2 instance that frequently polls the version control system todetect the new feature, use AWS CloudFormation and Amazon EC2 user data to runany testing harnesses to verify application functionality and then use AmazonSNS to notify the development team of the results.
D. Usean Elastic Beanstalk application that polls the version control system todetect the new feature, use AWS CloudFormation and Amazon EC2 user data to runany testing harnesses to verify application functionality and then use AmazonKinesis to notify the development team of the results.
E. UseAWS CloudFormation to launch an Amazon EC2 instance use Amazon EC2 user data torun any testing harnesses to verify application functionality and then useAmazon Kinesis to notify the development team of the results.

Answer: A

Explanation:
Using Amazon Kinesis would just take more time in setup and would not be ideal to notify the relevant team in the shortest time possible.
Since the test needs to be conducted in the staging VPC, it is best to launch the CC2 in the staging VPC.
For more information on the Simple Notification service, please visit the link:
• https://aws.amazon.com/sns/

NEW QUESTION 227
You have a web application composed of an Auto Scaling group of web servers behind a load balancer, and create a new AMI for each application version for deployment. You have a new version to release, and you want to use the A/B deployment technique to migrate users over in a controlled manner while the size of the fleet remains constant over a period of 12 hours, to ensure that the new version is performing well.
What option should you choose to enable this technique while being able to roll back easily?

A. Createan Auto scaling launch configuration with the new AM
B. Configure the AutoScalinggroup with the new launch configuratio
C. Use the Auto Scaling rollingupdates feature to migrate to the new version.
D. Createan Auto Scaling launch configuration with the new AM
E. Create an Auto Scalinggroup configured to use the new launch configuration and to register instanceswith the same load balance
F. Vary the desired capacity of each group tomigrate.
G. Createan Auto scaling launch configuration with the new AM
H. Configure Auto Scalingto vary the proportion of instances launched from the two launchconfigurations.
I. Createa load balance
J. Create an Auto Scaling launch configuration with the new AMIto use the new launch configuration and to registerinstances with the new loadbalance
K. Use Amazon Route53 weighted Round Robin to vary the proportion ofrequests sent to the load balancers.
L. Launchnew instances using the new AMI and attach them to the Auto Scalinggroup.Configure Elastic Load Balancing to vary the proportion of requests sent toinstances running the two application versions.

Answer: D

Explanation:
Since you want to control the usage to the new application in a controlled manner, the best way is to use Route53 weighted method. The AWS documentation mentions the following on this method
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 Weighted Round Robin method, please visit the link: http://docs^ws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html/rrouting-policy- weighted

NEW QUESTION 232
You have an I/O and network-intensive application running on multiple Amazon EC2 instances that cannot handle a large ongoing increase in traffic. The Amazon EC2 instances are using two Amazon EBS PIOPS volumes each, and each instance is identical.
Which of the following approaches should be taken in order to reduce load on the instances with the least disruption to the application?

A. Createan AMI from each instance, and set up Auto Scaling groups with a largerinstance type that has enhanced networking enabled and is Amazon EBS-optimized.
B. Stopeach instance and change each instance to a larger Amazon EC2 instance typethat has enhanced networking enabled and is Amazon EBS-optimize
C. Ensure thatRAID striping is also set up on each instance.
D. Addan instance-store volume for each running Amazon EC2 instance and implementRAID striping to improve I/O performance.
E. Addan Amazon EBS volume for each running Amazon EC2 instance and implement RAIDstripingto improve I/O performance.
F. Createan AMI from an instance, and set up an Auto Scaling group with an instance typethat has enhanced networking enabled and is Amazon EBS-optimized.

Answer: E

Explanation:
The AWS Documentation mentions the following on AMI's
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 AMI's, please visit the link:
• http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/AMIs.html

NEW QUESTION 235
Your public website uses a load balancer and an Auto Scalinggroup in a virtual private cloud. Your chief security officer has asked you to set up a monitoring system that quickly detects and alerts your team when a large sudden traffic increase occurs. How should you set this up?

A. Setup an Amazon CloudWatch alarm for the Elastic Load Balancing Networkln metricand then use Amazon SNS to alert your team.
B. Usean Amazon EMR job to run every thirty minutes, analyze the Elastic LoadBalancing access logs in a batch manner to detect a sharp increase in trafficand then use the Amazon Simple Email Service to alert your team.
C. Usean Amazon EMR job to run every thirty minutes analyze the CloudWatch logs fromyour application Amazon EC2 instances in a batch manner to detect a sharpincrease in traffic and then use the Amazon SNS SMS notification to alert yourteam

D. Setup an Amazon CloudWatch alarm for the Amazon EC2 Networkln metric for the AutoScaling group and then use Amazon SNS to alert your team.
E. Setup a cron job to actively monitor the AWS CloudTrail logs for increased trafficand use Amazon SNS to alert your team.

Answer: D

Explanation:
The below snapshot from the AWS documentation gives details on the Networkln metric

NetworkIn	The number of bytes received on all network interfaces by the instance. This metric identifies the volume of incoming network traffic to a single instance. The number reported is the number of bytes received during the period. If you are using basic (five-minute) monitoring, you can divide this number by 300 to find Bytes/second. If you have detailed (one-minute) monitoring, divide it by 60. Units: Bytes

NEW QUESTION 238
As part of your deployment process, you are configuring your continuous integration (CI) system to build AMIs. You want to build them in an automated manner that is also cost-efficient. Which method should you use?

A. Attachan Amazon EBS volume to your CI instance, build the root file system of yourimage on the volume, and use the CreateImage API call to create an AMI out ofthis volume.
B. Havethe CI system launch a new instance, bootstrap the code and apps onto theinstance and create an AMI out of it.
C. Uploadall contents of the image to Amazon S3 launch the base instance, download allof the contents from Amazon S3 and create the AMI.
D. Havethe CI system launch a new spot instance bootstrap the code and apps onto theinstance and create an AMI out of it.

Answer: D

Explanation:
The AWS documentation mentions the following
If your organization uses Jenkins software in a CI/CD pipeline, you can add Automation as a post- build step to pre-install application releases into Amazon Machine Images (AMIs). You can also use the Jenkins scheduling feature to call Automation and create your own operating system (OS) patching cadence
For more information on Automation with Jenkins, please visit the link:
• http://docs.aws.a mazon.com/systems-manager/latest/userguide/automation-jenkinsJntm I
• https://wiki.jenkins.io/display/JCNKINS/Amazon < CC21 Plugin

NEW QUESTION 240
As part of your deployment pipeline, you want to enable automated testing of your AWS CloudFormation template. What testing should be performed to enable faster feedback while minimizing costs and risk? Select three answers from the options given below

A. Usethe AWS CloudFormation Validate Template to validate the syntax of the template
B. Usethe AWS CloudFormation Validate Template to validate the properties ofresources defined in the template.
C. Validatethe template's is syntax using a generalJSON parser.
D. Validatethe AWS CloudFormation template against the official XSD scheme definitionpublished by Amazon Web Services.
E. Updatethe stack with the templat
F. If the template fails rollback will return thestack and its resources to exactly the same state.
G. When creating the stack, specify an Amazon SNS topic to which your testing system is subscribe
H. Your testing system runs tests when it receives notification that the stack is created or updated.

Answer: AEF

Explanation:
The AWS documentation mentions the following
The aws cloudformation validate-template command is designed to check only the syntax of your template. It does not ensure that the property values that you have specified for a resource are valid for that resource. Nor does it determine the number of resources that will exist when the stack is created.
To check the operational validity, you need to attempt to create the stack. There is no sandbox or test area for AWS Cloud Formation stacks, so you are charged for the resources you create during testing. Option F is needed for notification.
For more information on Cloudformation template validation, please visit the link:
http://docs.aws.a mazon.com/AWSCIoudFormation/latest/UserGuide/using-cfn-va I idate- template.htm I

NEW QUESTION 244
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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