Exam Questions DOP-C01

AWS Certified DevOps Engineer- Professional

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

NEW QUESTION 1
You have decided that you need to change the instance type of your production instances which are running as part of an AutoScaling group. The entire architecture is deployed using CloudFormation Template. You currently have 4 instances in Production. You cannot have any interruption in service and need to ensure 2 instances are always runningduring the update? Which of the options below listed can be used for this?

A. AutoScalingRollingUpdate
B. AutoScalingScheduledAction
C. AutoScalingReplacingUpdate
D. AutoScalingIntegrationUpdate

Answer: A

Explanation:
The AWS::AutoScaling::AutoScalingGroup resource supports an UpdatePolicy attribute. This is used to define how an Auto Scalinggroup resource is updated when an update to the Cloud Formation stack occurs. A common approach to updating an Auto Scaling group is to perform a rolling update, which is done by specifying the AutoScalingRollingUpdate policy. This retains the same Auto Scaling group and replaces old instances with new ones, according to the parameters specified. For more information on Autoscaling updates, please refer to the below link: https://aws.amazon.com/premiumsupport/knowledge-center/auto-scaling-group-rolling-updates/

NEW QUESTION 2
You have an ELB setup in AWS with EC2 instances running behind it. You have been requested to monitor the incoming connections to the ELB. Which of the below options can suffice this requirement?

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

Answer: B

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

NEW QUESTION 3
You are responsible for your company's large multi-tiered Windows-based web application running on Amazon EC2 instances situated behind a load balancer. While reviewing metrics, you've started noticing an upwards trend for slow customer page load time. Your manager has asked you to come up with a solution to ensure that customer load time is not affected by too many requests per second. Which technique would you use to solve this issue?

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

Answer: C

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

NEW QUESTION 4
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 5
You currently run your infrastructure on Amazon EC2 instances behind an Auto Scalinggroup. All logs for your application are currentl\ written to ephemeral storage. Recently your company experienced a major bug in the code that made it through testing and was ultimately deployed to your fleet. This bug triggered your Auto Scalinggroup to scale up and back down before you could successfully retrieve the logs off your server to better assist you in troubleshooting the bug. Which technique should you use to make sure you are able to review your logs after your instances have shut down?

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

Answer: C

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

NEW QUESTION 6
You have the following application to be setup in AWS
1) A web tier hosted on EC2 Instances
2) Session data to be written to DynamoDB
3) Log files to be written to Microsoft SQL Server
How can you allow an application to write data to a DynamoDB table?

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

Answer: D

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

NEW QUESTION 7
You have an Auto Scaling group of Instances that processes messages from an Amazon Simple Queue Service (SQS) queue. The group scales on the size of the queue. Processing Involves calling a third- party web service. The web service is complaining about the number of failed and repeated calls it is receiving from you. You have noticed that when the group scales in, instances are being terminated while they are processing. What cost-effective solution can you use to reduce the number of incomplete process attempts?

A. Create a new Auto Scaling group with minimum and maximum of 2 and instances running web proxy softwar
B. Configure the VPC route table to route HTTP traffic to these web proxies.
C. Modify the application running on the instances to enable termination protection while it processes a task and disable it when the processing is complete.
D. Increase the minimum and maximum size for the Auto Scalinggroup, and change the scaling policies so they scale less dynamically.
E. Modify the application running on the instances to put itself into an Auto Scaling Standby state while it processes a task and return itself to InService when the processing is complete.

Answer: D

Explanation:
The following diagram shows the lifecycle of the instances in Autoscaling

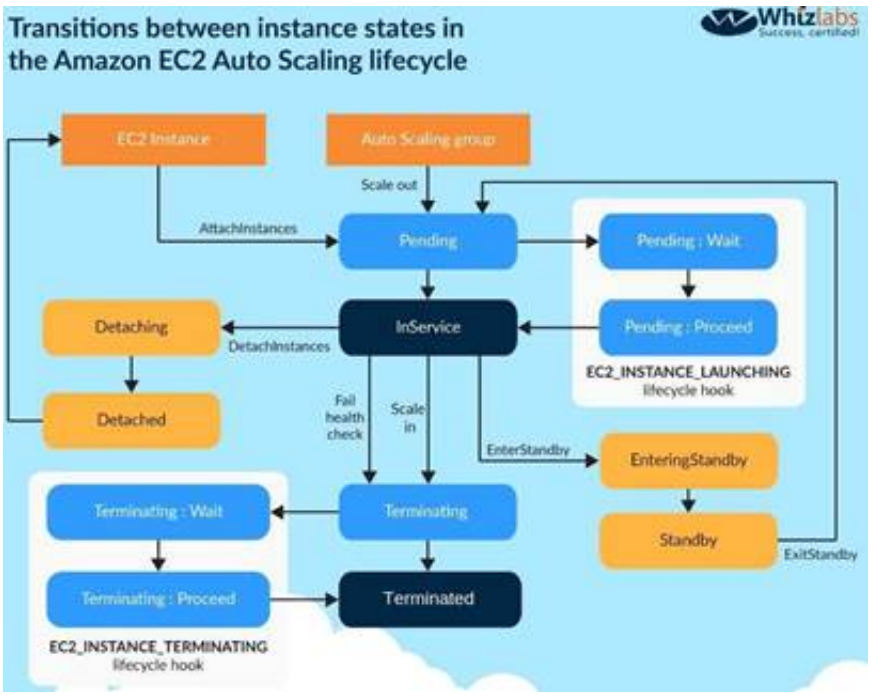

You can put the instances in a standby state, via the application, do the processing and then put the instance back in a state where it can be governed by the Autoscaling Group.
For more information on the Autoscaling Group Lifecycle please refer to the below link:
http://docs.aws.amazon.com/autoscaling/latest/userguide/AutoScaingGroupl_ifecycle.htm I Note: As per AWS documentation.
To control whether an Auto Scaling group can terminate a particular instance when scaling in, use instance protection.
It is termed as Instance protection rather than termination protection when we refer it with "Scaling in process" of ASG.
For more information please view the following link: https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-instance-termination.htmlffinstance- protection-instance

NEW QUESTION 8
Your mobile application includes a photo-sharing service that is expecting tens of thousands of users at launch. You will leverage Amazon Simple Storage Service (S3) for storage of the user Images, and you must decide how to authenticate and authorize your users for access to these images. You also need to manage the storage of these images. Which two of the following approaches should you use? Choose two answers from the options below

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

Answer: CE

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

NEW QUESTION 9
After reviewing the last quarter's monthly bills, management has noticed an increase in the overall bill from Amazon. After researching this increase in cost, you discovered that one of your new services is doing a lot of GET Bucket API calls to Amazon S3 to build a metadata cache of all objects in the applications bucket. Your boss has asked you to come up with a new cost-effective way to help reduce the amount of these new GET Bucket API calls. What process should you use to help mitigate the cost?

A. Update your Amazon S3 buckets' lifecycle policies to automatically push a list of objects to a new bucket, and use this list to view objects associated with the application's bucket.
B. Create a new DynamoDB tabl
C. Use the new DynamoDB table to store all metadata about all objects uploaded to Amazon S3. Any time a new object is uploaded, update the application's internalAmazon S3 object metadata cache from DynamoDB.C Using Amazon SNS, create a notification on any new Amazon S3 objects that automatical ly updates a new DynamoDB table to store allmetadata about the new objec
D. Subscribe the application to the Amazon SNS topic to update its internal Amazon S3 object metadata cache from the DynamoDB tabl
E. ^/
F. Upload all files to an ElastiCache file cache serve
G. Update your application to now read all file metadata from the ElastiCache file cache server, and configure the ElastiCache policies to push all files to Amazon S3 for long-term storage.

Answer: C

Explanation:
Option A is an invalid option since Lifecycle policies are normally used for expiration of objects or archival of objects.
Option B is partially correct where you store the data in DynamoDB, but then the number of GET requests would still be high if the entire DynamoDB table had to be
traversed and each object compared and updated in S3.
Option D is invalid because uploading all files to Clastic Cache is not an ideal solution.
The best option is to have a notification which can then trigger an update to the application to update the DynamoDB table accordingly.
For more information on SNS triggers and DynamoDB please refer to the below link:
? https://aws.amazon.com/blogs/compute/619/

NEW QUESTION 10
You have a complex system that involves networking, 1AM policies, and multiple, three-tier applications. You are still receiving requirements for the new system, so you don't yet know how many AWS components will be present in the final design. You want to start using AWS CloudFormation to define these AWS resources so that you can automate and version-control your infrastructure. How would you use AWS CloudFormation to provide agile new environments for your customers in a cost-effective, reliable manner?

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

Answer: B

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

NEW QUESTION 10
One of the instances in your Auto Scaling group health check returns the status of Impaired to Auto Scaling. What will Auto Scaling do in this case.

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

Answer: A

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

NEW QUESTION 12
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 13
As an architect you have decided to use CloudFormation instead of OpsWorks or Elastic Beanstalk for deploying the applications in your company. Unfortunately, you have discovered that there is a
resource type that is not supported by CloudFormation. What can you do to get around this.

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

Answer: B

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

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

NEW QUESTION 17
You have deployed an Elastic Beanstalk application in a new environment and want to save the current state of your environment in a document. You want to be able to restore your environment to the current state later or possibly create a new environment. You also want to make sure you have a restore point. How can you achieve this?

A. Use CloudFormation templates
B. Configuration Management Templates
C. Saved Configurations
D. Saved Templates

Answer: C

Explanation:
You can save your environment's configuration as an object in Amazon S3 that can be applied to other environments during environment creation, or applied to a running environment. Saved configurations are YAML formatted templates that define an environment's platform configuration, tier, configuration option settings, and tags.
For more information on Saved Configurations please refer to the below link:
• http://docs.aws.a mazon.com/elasticbeanstalk/latest/dg/envi ronment-configuration- savedconfig.html

NEW QUESTION 22
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.

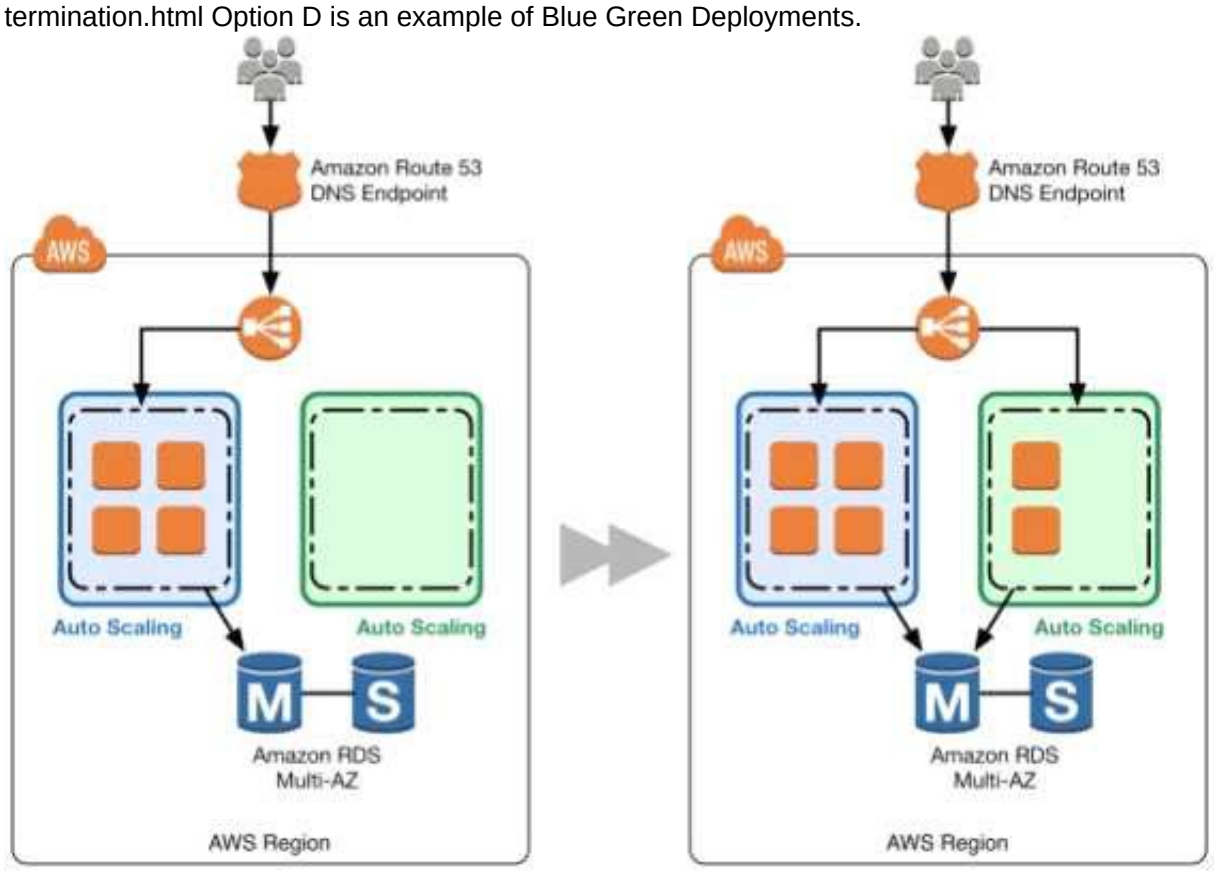

A blue group carries the production load while a green group is staged and deployed with the new code. When if s time to deploy, you simply attach the green group to the existing load balancer to introduce traffic to the new environment. For HTTP/HTTP'S listeners, the load balancer favors the green Auto 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 25
You are using a configuration management system to manage your Amazon EC2 instances. On your Amazon EC2 Instances, you want to store credentials for connecting to an Amazon RDS MYSQL DB instance. How should you securely store these credentials?

A. Give the Amazon EC2 instances an 1AM role that allows read access to a private Amazon S3 bucke
B. Store a file with database credentials in the Amazon S3 bucke

C. Have your configuration management system pull the file from the bucket when it is needed.
D. Launch an Amazon EC2 instance and use the configuration management system to bootstrap the instance with the Amazon RDS DB credential
E. Create an AMI from this instance.
F. Store the Amazon RDS DB credentials in Amazon EC2 user dat
G. Import the credentials into the Instance on boot.
H. Assign an 1AM role to your Amazon EC2 instance, and use this 1AM role to access the Amazon RDS DB from your Amazon EC2 instances.

Answer: D

Explanation:
Creating and Using an 1AM Policy for 1AM Database Access
To allow an 1AM user or role to connect to your DB instance or DB cluster, you must create an 1AM policy. After that you attach the policy to an 1AM user or role.
Note
To learn more about 1AM policies, see Authentication and Access Control for Amazon RDS.
The following example policy allows an 1AM user to connect to a DB instance using 1AM database authentication.

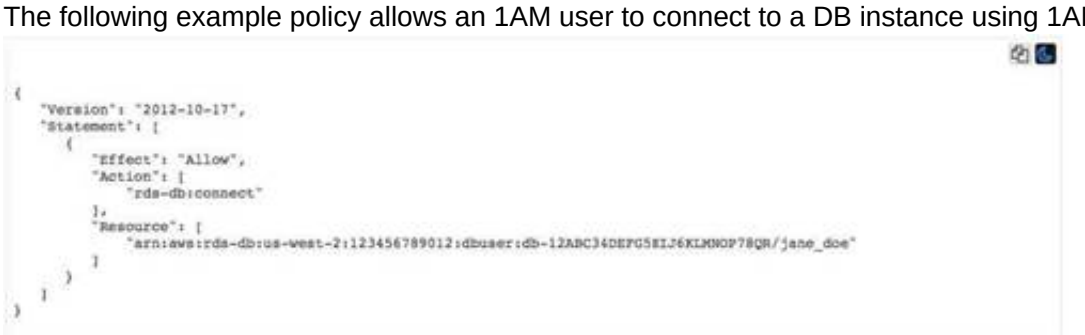

```
{
   "Version": "2012-10-17",
   "Statement": [
      {
         "Effect": "Allow",
         "Action": [
            "rds-db:connect"
         ],
         "Resource": [
            "arn:aws:rds-db:us-west-2:123456789012:dbuser:db-12ABC34DEFG5HIJ6KLMNOP78QR/jane_doe"
         ]
      }
   ]
}
```

Important
Don't confuse the rds-db: prefix with other Amazon RDS action prefixes that begin with rds:. You use the rds-db: prefix and the rds-db:connect action only for 1AM database authentication. They aren't valid in any other context.
1AM Database Authentication for MySQL and Amazon Aurora
With Amazon RDS for MySQL or Aurora with MySQL compatibility, you can authenticate to your DB instance or DB cluster using AWS Identity and Access Management (IAMJ database authentication. With this authentication method, you don't need to use a password when you connect to a DB instance. Instead, you use an authentication token.
An authentication token is a unique string of characters that Amazon RDS generates on request. Authentication tokens are generated using AWS Signature Version 4. Each token has a lifetime of 15 minutes. You don't need to store user credentials in the database, because authentication is managed externally using 1AM. You can also still use standard database authentication.
IAM database authentication provides the following benefits:
• Network traffic to and from the database is encrypted using Secure Sockets Layer (SSL).
• You can use IAM to centrally manage access to your database resources, instead of managing access individually on each DB instance or DB cluster.
• For applications running on Amazon EC2, you can use EC2 instance profile credentials to access the database instead of a password, for greater security.
For more information please refer to the below document link from AWS
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.IAMDBAuth.html
https://docs^ws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.IAMDBAuth.IAMPolicy. html
You can use roles to delegate access to users, applications, or services that don't normally have access to your AWS resources. For example, you might want to grant users in your AWS account access to resources they don't usually have, or grant users in one AWS account access to resources in another account. Or you might want to allow a mobile app to use AWS resources, but not want to embed AWS keys within the app (where they can be difficult to rotate and where users can potentially extract them). Sometimes you want to give AWS access to users who already have identities defined outside of AWS, such as in your corporate directory. Or, you might want to grant access to your account to third parties so that they can perform an audit on your resources. For more information on 1AM Roles, please refer to the below document link: from AWS
http://docs.aws.a mazon.com/IAM/latest/UserGuide/id_roles. html

NEW QUESTION 26
You have an application hosted in AWS. This application was created using Cloudformation Templates and Autoscaling. Now your application has got a surge of users which is decreasing the performance of the application. As per your analysis, a change in the instance type to C3 would resolve the issue. Which of the below option can introduce this change while minimizing downtime for end users?

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

Answer: B

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

NEW QUESTION 28

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 32
There is a requirement to monitor API calls against your AWS account by different users and entities. There needs to be a history of those calls. The history of those calls are needed in in bulk for later review. Which 2 services can be used in this scenario

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

Answer: C

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

NEW QUESTION 33
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 37
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 39
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 42
You need to create a simple, holistic check for your system's general availablity and uptime. Your system presents itself as an HTTP-speaking API. What is the most simple tool on AWS to achieve this with?

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

Answer: A

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

NEW QUESTION 45
You need to scale an RDS deployment. You are operating at 10% writes and 90% reads, based on your logging. How best can you scale this in a simple way?

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

Answer: C

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

NEW QUESTION 48
Your company needs to automate 3 layers of a large cloud deployment. You want to be able to track this deployment's evolution as it changes over time, and carefully control any alterations. What is a good way to automate a stack to meet these requirements?

A. Use OpsWorks Stacks with three layers to model the layering in your stack.
B. Use CloudFormation Nested Stack Templates, with three child stacks to represent the three logical layers of your cloud.

C. Use AWS Config to declare a configuration set that AWS should roll out to your cloud.
D. Use Elastic Beanstalk Linked Applications, passing the important DNS entires between layers using the metadata interface.

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 nested stacks, please visit the below URL:
• http://docs^ws.amazon.com/AWSCIoudFormation/latest/UserGuide/best-practices.html#nested Note:
The query is, how you can automate a stack over the period of time, when changes are required, with out recreating the stack.
The function of Nested Stacks are to reuse Common Template Patterns.
For example, assume that you have a load balancer configuration that you use for most of your stacks. Instead of copying and pasting the same configurations into your templates, you can create a dedicated template for the load balancer. Then, you just use the resource to reference that template from within other templates.
Yet another example is if you have a launch configuration with certain specific configuration and you need to change the instance size only in the production environment and to leave it as it is in the development environment.
AWS also recommends that updates to nested stacks are run from the parent stack.
When you apply template changes to update a top-level stack, AWS CloudFormation updates the top-level stack and initiates an update to its nested stacks. AWS Cloud Formation updates the resources of modified nested stacks, but does not update the resources of unmodified nested stacks.

NEW QUESTION 50
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 51
You need to run a very large batch data processingjob one time per day. The source data exists
entirely in S3, and the output of the processingjob should also be written to S3 when finished. If you need to version control this processingjob and all setup and teardown logic for the system, what approach should you use?.

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

Answer: B

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

NEW QUESTION 52
There is a very serious outage at AWS. EC2 is not affected, but your EC2 instance deployment scripts
stopped working in the region with the outage. What might be the issue?

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

Answer: B

Explanation:

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

NEW QUESTION 53
You need your API backed by DynamoDB to stay online duringa total regional AWS failure. You can tolerate a couple minutes of lag or slowness during a large failure event, but the system should recover with normal operation after those few minutes. What is a good approach?

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

Answer: B

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

NEW QUESTION 55
You need to create an audit log of all changes to customer banking data. You use DynamoDB to store this customer banking data. It's important not to lose any information due to server failures. What is an elegant way to accomplish this?

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

Answer: A

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

NEW QUESTION 59
You are building a Ruby on Rails application for internal, non-production use which uses MySQL as a database. You want developers without very much AWS experience to be able to deploy new code with a single command line push. You also want to set this up as simply as possible. Which tool is ideal for this setup?

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

Answer: D

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

NEW QUESTION 60

You run a 2000-engineer organization. You are about to begin using AWS at a large scale for the first time. You want to integrate with your existing identity management system running on Microsoft Active Directory, because your organization is a power-user of Active Directory. How should you manage your AWS identities in the most simple manner?

A. Use AWS Directory Sen/ice Simple AD.
B. Use AWS Directory Service AD Connector.
C. Use an Sync Domain running on AWS Directory Sen/ice.
D. Use an AWS Directory Sync Domain running on AWS Lambda.

Answer: B

Explanation:
AD Connector is a directory gateway with which you can redirect directory requests to your on- premises Microsoft Active Directory without caching any information in the cloud. AD Connector comes in two sizes, small and large. A small AD Connector is designed for smaller organizations of up to 500 users. A large AD Connector
can support larger organizations of up to 5,000 users. Once set up, AD Connector offers the following benefits:
• Your end users and IT administrators can use their existing corporate credentials to log on to AWS applications such as Amazon Workspaces, Amazon WorkDocs, or Amazon WorkMail.
• You can manage AWS resources like Amazon EC2 instances or Amazon S3 buckets through 1AM role-based access to the AWS Management Console.
• You can consistently enforce existing security policies (such as password expiration, password history, and account lockouts) whether users or IT administrators are accessing resources in your on- premises infrastructure or in the AWS Cloud.
• You can use AD Connector to enable multi-factor authentication by integrating with your existing RADIUS-based MFA infrastructure to provide an additional layer of security when users access AWS applications.
For more information on the AD Connector, please visit the below URL:
• http://docs.aws.amazon.com/directoryservice/latest/admin-guide/directory_ad_con nector.htm l

NEW QUESTION 61
You need to deploy an AWS stack in a repeatable manner across multiple environments. You have selected CloudFormation as the right tool to accomplish this, but have found that there is a resource type you need to create and model, but is unsupported by CloudFormation. How should you overcome this challenge?

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

Answer: D

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

NEW QUESTION 64
Your API requires the ability to stay online during AWS regional failures. Your API does not store any state, it only aggregates data from other sources - you do not have a database. What is a simple but effective way to achieve this uptime goal?

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

Answer: D

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

NEW QUESTION 67
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 72
Your development team is using access keys to develop an application that has access to S3 and DynamoDB. A new security policy has outlined that the credentials should not be older than 2 months, and should be rotated. How can you achieve this

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

Answer: B

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

NEW QUESTION 76
An EC2 instance has failed a health check. What will the ELB do?

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

Answer: B

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

NEW QUESTION 77
You have carried out a deployment using Elastic Beanstalk with All at once method, but the application is unavailable. What could be the reason for this

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

Answer: C

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

NEW QUESTION 81
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 85
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 87
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 91
Which of the below services can be used to deploy application code content stored in Amazon S3 buckets, GitHub repositories, or Bitbucket repositories

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

Answer: B

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

NEW QUESTION 96
Which of the following Deployment types are available in the CodeDeploy service. Choose 2 answers from the options given below

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

Answer: AD

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

NEW QUESTION 101

When your application is loaded onto an Opsworks stack, which of the following event is triggered by Opsworks?

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

Answer: A

Explanation:
When you deploy an application, AWS Ops Works Stacks triggers a Deploy event, which runs each layer's Deploy recipes. AWS OpsWorks Stacks also installs stack configuration and deployment attributes that contain all of the information needed to deploy the app, such as the app's repository and database connection data. For more information on the Deploy event please refer to the below link:
• http://docs.aws.amazon.com/opsworks/latest/userguide/workingapps.html

NEW QUESTION 102
Which of the following Cache Engines does Opswork have built in support for?

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

Answer: B

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

NEW QUESTION 106
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 109
Which of the following is not a rolling type update which is present for Configuration Updates when it comes to the Elastic Beanstalk service

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

Answer: B

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

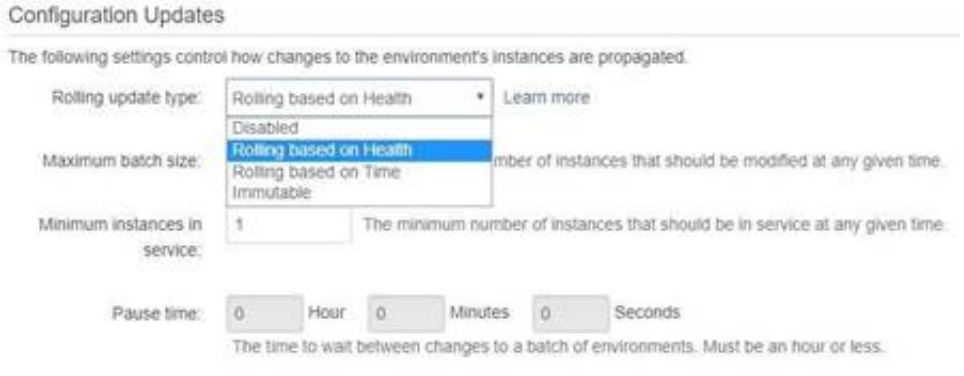

The AWS Documentation mentions
1) With health-based rolling updates. Elastic Beanstalk waits until instances in a batch pass health checks before moving on to the next batch.

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

NEW QUESTION 113
Your company has a set of resources hosted in AWS. Your IT Supervisor is concerned with the costs being incurred by the resources running in AWS and wants to optimize on the costs as much as possible. Which of the following ways could help achieve this efficiently? Choose 2 answers from the options given below.

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

Answer: AB

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

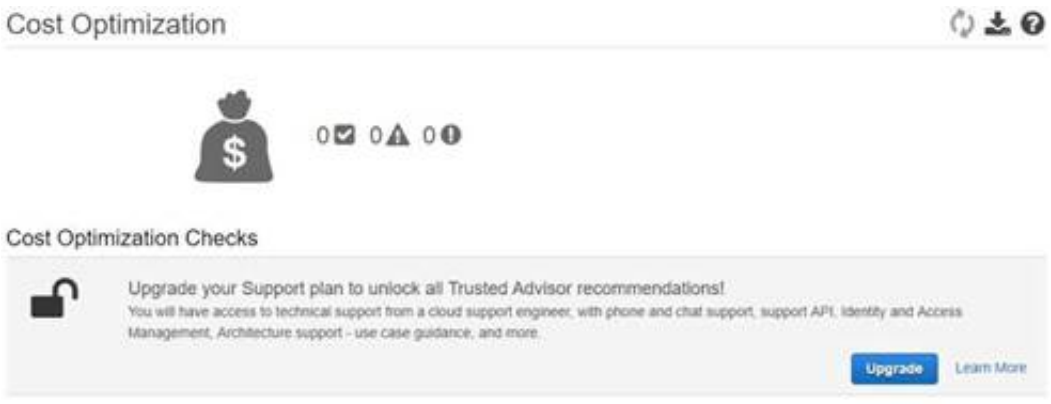

For more information on the Trusted Advisor, please visit the below U RL:
• https://aws.amazon.com/premiumsupport/trustedadvisor/

NEW QUESTION 117
Your application is having a very high traffic, so you have enabled autoscaling in multi availability zone to suffice the needs of your application but you observe that one of the availability zone is not receiving any traffic. What can be wrong here?

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

Answer: C

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

NEW QUESTION 119
You have a legacy application running that uses an m4.large instance size and cannot scale with Auto Scaling, but only has peak performance 5% of the time. This is a huge waste of resources and money so your Senior Technical Manager has set you the task of trying to reduce costs while still keeping the legacy application running as it should. Which of the following would best accomplish the task your manager has set you? Choose the correct answer from the options below

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

Answer: A

Explanation:
The aws documentation clearly indicates using T2 CC2 instance types for those instances which don't use CPU that often.
T2
T2 instances are Burstable Performance Instances that provide a baseline level of CPU performance with the ability to burst above the baseline.
T2 Unlimited instances can sustain high CPU performance for as long as a workload needs it. For most general-purpose workloads, T2 Unlimited instances will provide ample performance without any additional charges. If the instance needs to run at higher CPU utilization for a prolonged period, it can also do so at a flat additional charge of 5 cents per vCPU-hour.
The baseline performance and ability to burst are governed by CPU Credits. T2 instances receive CPU Credits continuously at a set rate depending on the instance size, accumulating CPU Credits when they are idle, and consuming CPU credits when they are active. T2 instances are a good choice for a variety of

general-purpose workloads including micro-services, low-latency interactive applications, small and medium databases, virtual desktops, development, build and stage environments, code repositories, and product prototypes. For more information see Burstable Performance Instances.
For more information on F_C2 instance types please see the below link: https://aws.amazon.com/ec2/instance-types/

NEW QUESTION 122
Your security officer has told you that you need to tighten up the logging of all events that occur on your AWS account. He wants to be able to access all events that occur on the account across all regions quickly and in the simplest way possible. He also wants to make sure he is the only person that has access to these events in the most secure way possible. Which of the following would be the best solution to assure his requirements are met? Choose the correct answer from the options below

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

Answer: A

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

NEW QUESTION 125
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 130
An enterprise wants to use a third-party SaaS application running on AWS.. The SaaS application needs to have access to issue several API commands to discover Amazon EC2 resources running within the enterprise's account. The enterprise has internal security policies that require any outside access to their environment must conform to the principles of least privilege and there must be controls in place to ensure that the credentials used by the SaaS vendor cannot be used by any other third party. Which of the following would meet all of these conditions?

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

Answer: C

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

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

NEW QUESTION 135
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 140
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 143
You are planning on using AWS Code Deploy in your AWS environment. Which of the below features of AWS Code Deploy can be used to Specify scripts to be run on each instance at various stages of the deployment process

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

Answer: A

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

NEW QUESTION 148
You work for a company that automatically tags photographs using artificial neural networks (ANNs), which run on GPUs using C++. You receive millions of images at a time, but only 3 times per day on average. These images are loaded into an AWS S3 bucket you control for you in a batch, and then the customer publishes a JSON-formatted manifest into another S3 bucket you control as well. Each image takes 10 milliseconds to process using a full GPU. Your neural network software requires 5 minutes to bootstrap. Image tags are JSON objects, and you must publish them to an S3 bucket. Which of these is the best system architectures for this system?

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

Answer: B

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

NEW QUESTION 150
You are currently using SGS to pass messages to EC2 Instances. You need to pass messages which are greater than 5 MB in size. Which of the following can help you accomplish this.

A. UseKinesis as a buffer stream for message bodie
B. Store the checkpoint id fortheplacement in the Kinesis Stream in SQS.
C. Usethe Amazon SQS Extended Client Library for Java and Amazon S3 as a storagemechanism for message bodie
D. */
E. UseSQS's support for message partitioning and multi-part uploads on Amazon S3.
F. UseAWS EFS as a shared pool storage mediu
G. Store filesystem pointers to the fileson disk in the SQS message bodies.

Answer: B

Explanation:
The AWS documentation mentions the following
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 SQS and sending larger messages please visit the link

NEW QUESTION 154
A user is using Cloudformation to launch an EC2 instance and then configure an application after the instance is launched. The user wants the stack creation of ELB and AutoScaling to wait until the EC2 instance is launched and configured properly. How can the user configure this?

A. It is not possible that the stackcreation will wait until one service is created and launchedB.The user can use theHoldCondition resource to wait for the creation of the other dependent resources
B. The user can use theDependentCondition resource to hold the creation of the other dependentresources
C. The user can use the WaitConditionresource to hold the creation of the other dependent resources

Answer: D

Explanation:
You can use a wait condition for situations like the following:
To coordinate stack resource creation with configuration actions that are external to the stack creation
To track the status of a configuration process
For more information on Cloudformation Wait condition please visit the link
http://docs^ws.amazon.com/AWSCIoudFormation/latest/UserGuide/aws-properties-waitcondition.htmI

NEW QUESTION 158
A gaming company adopted AWS Cloud Formation to automate load-testing of theirgames. They have created an AWS Cloud Formation template for each gaming environment and one for the load- testing stack. The load-testing stack creates an Amazon Relational Database Service (RDS) Postgres database and two web servers running on Amazon Elastic Compute Cloud (EC2) that send HTTP requests, measure response times, and write the results into the database. A test run usually takes between 15 and 30 minutes. Once the tests are done, the AWS Cloud Formation stacks are torn down immediately. The test results written to the Amazon RDS database must remain accessible for visualization and analysis.
Select possible solutions that allow access to the test results after the AWS Cloud Formation load - testing stack is deleted.
Choose 2 answers.

A. Define an Amazon RDS Read-Replica in theload-testing AWS Cloud Formation stack and define a dependency relation betweenmaster and replica via the Depends On attribute.
B. Define an update policy to prevent deletionof the Amazon RDS database after the AWS Cloud Formation stack is deleted.
C. Define a deletion policy of type Retain forthe Amazon RDS resource to assure that the RDS database is not deleted with theAWS Cloud Formation stack.
D. Define a deletion policy of type Snapshotfor the Amazon RDS resource to assure that the RDS database can be restoredafter the AWS Cloud Formation stack is deleted.
E. Defineautomated backups with a backup retention period of 30 days for the Amazon RDSdatabase and perform point-in-time recovery of the database after the AWS CloudFormation stack is deleted.

Answer: CD

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

NEW QUESTION 162
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/AWSCIoudFormation/latest/UserGuide/intrinsic-function -reference, html

NEW QUESTION 166
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 167
You need to store a large volume of data. The data needs to be readily accessible for a short period, but then needs to be archived indefinitely after that. What is a cost-effective solution?

A. Storeall the data in S3 so that it can be more cost effective

B. Storeyour data in Amazon S3, and use lifecycle policies to archive to Amazon Glacier
C. Storeyour data in an EBS volume, and use lifecycle policies to archive to AmazonGlacier.
D. Storeyour data in Amazon S3, and use lifecycle policies to archive toS3-Infrequently Access

Answer: B

Explanation:
The AWS documentation mentions the following on Lifecycle policies
Lifecycle configuration enables you to specify the lifecycle management of objects in a bucket. The configuration is a set of one or more rules, where each rule defines an action for Amazon S3 to apply to a group of objects. These actions can be classified as follows:
Transition actions - In which you define when objects transition to another storage class. For example, you may choose to transition objects to the STANDARDJ A (IA, for infrequent access) storage class 30 days after creation, or archive objects to the GLACIER storage class one year after creation.
Expiration actions - In which you specify when the objects expire. Then Amazon S3 deletes the expired objects on your behalf. For more information on S3 Lifecycle policies, please visit the below URL
• http://docs.aws.a mazon.com/AmazonS3/latest/dev/object-lifecycle-mgmt.html

NEW QUESTION 170
Explain what the following resource in a CloudFormation template does? Choose the best possible answer.

```
"SNSTopic" : {
    "Type" : "AWS::SNS::Topic",
    "Properties" : {
        "Subscription" : [{
            "Protocol" : "sqs",
            "Endpoint" : { "Fn::GetAtt" : [ "SQSQueue", "Arn" ] }
        }]
    }
```

A. Createsan SNS topic which allows SQS subscription endpoints to be added as a parameteron thetemplate
B. Createsan SNS topic that allow SQS subscription endpoints
C. Createsan SNS topic and then invokes the call to create an SQS queue with a logicalresource name of SQSQueue
D. Creates an SNS topic and adds asubscription ARN endpoint for the SQS resource created under the logical nameSQSQueue

Answer: D

Explanation:
The intrinsic function Fn::GetAtt returns the value of an attribute from a resource in the template. This has nothing to do with adding parameters (Option A is wrong) or allowing endpoints (Option B is wrong) or invoking relevant calls (Option C is wrong)
For more information on Fn:: GetAtt function please refer to the below link
http://docs.aws.a mazon.com/AWSCIoudFormation/latest/UserGuide/intrinsic-function -reference- getatt.htm I

NEW QUESTION 172
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 174
You have a number of Cloudformation stacks in your IT organization. Which of the following commands will help see all the cloudformation stacks which have a completed status?

A. describe-stacks
B. list-stacks
C. stacks-complete
D. list-templates

Answer: B

Explanation:
The following is the description of the list-stacks command
Returns the summary information for stacks whose status matches the specified StackStatusFilter.
Summary information for stacks that have been deleted is kept for 90 days after the stack is deleted. If no stack-status-filter is specified, summary information for all stacks is returned (including existing stacks and stacks that have been deleted).
For more information on the list-stacks command please visit the below link http://docs.aws.amazon.com/cli/latest/reference/cloudformation/list-stacks. html

NEW QUESTION 177
If you're trying to configure an AWS Elastic Beanstalk worker tier for easy debugging if there are problems finishing queue jobs, what should you configure?

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

Answer: D

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

NEW QUESTION 181
You are in charge of designing a Cloudformation template which deploys a LAMP stack. After deploying a stack, you see that the status of the stack is showing as CREATE_COMPLETE, but the apache server is still not up and running and is experiencing issues while starting up. You want to ensure that the stack creation only shows the status of CREATE_COMPLETE after all resources defined in the stack are up and running. How can you achieve this?
Choose 2 answers from the options given below.

A. Definea stack policy which defines that all underlying resources should be up andrunning before showing a status of CREATE_COMPLETE.
B. Uselifecycle hooks to mark the completion of the creation and configuration of theunderlying resource.
C. Usethe CreationPolicy to ensure it is associated with the EC2 Instance resource.
D. Usethe CFN helper scripts to signal once the resource configuration is complete.

Answer: CD

Explanation:
The AWS Documentation mentions
When you provision an Amazon EC2 instance in an AWS Cloud Formation stack, you might specify additional actions to configure the instance, such as install software packages or bootstrap applications. Normally, CloudFormation proceeds with stack creation after the instance has been successfully created. However, you can use a Creation Pol icy so that CloudFormation proceeds with stack creation only after your configuration actions are done. That way you'll know your applications are ready to go after stack creation succeeds.
For more information on the Creation Policy, please visit the below url https://aws.amazon.com/blogs/devops/use-a-creationpolicy-to-wait-for-on-instance-configurations/

NEW QUESTION 185
Which of the following is false when it comes to using the Elastic Load balancer with Opsworks stacks?

A. Youcan attach only one load balancer to a layer.
B. A Classic Load Balancer can span across AWSOpsWorks Stacks layers.
C. Eachload balancer can handle only one layer.
D. Youneed to create the load balancer before hand and then attach it to the Opsworkstack.

Answer: B

Explanation:
The AWS Documentation mentions the following
To use Clastic Load Balancing with a stack, you must first create one or more load balancers in the same region by using the Clastic Load Balancing console, CLI, or API. You should be aware of the following:
You can attach only one load balancer to a layer. Cach load balancer can handle only one layer.
AWS OpsWorks Stacks does not support Application Load Balancer. You can only use Classic Load Balancer with AWS OpsWorks Stacks. For more information on Clastic Load Balancer with Opswork,
please visit the below url http://docs.aws.a mazon.com/opsworks/latest/userguide/layers-elb.html

NEW QUESTION 189
You want to use Code Deploy to deploy code that is hosted on yourgithub repository. Which of the following additional services can help fulfil this requirement.

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

Answer: A

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

NEW QUESTION 194
Which of the following tools for EC2 can be used to administer instances without the need to SSH or RDP into the instance.

A. AWSConfig
B. AWSCodePipeline
C. RunCommand
D. EC2Config

Answer: C

Explanation:
You can use Run Command from the Amazon L~C2 console to configure instances without having to login to each instance
For more information on the Run Command, please visit the below URL:
• http://docs.aws.a mazon.com/systems-manager/latest/userguide/rc-console.html

NEW QUESTION 196
You are trying to debug the creation of Cloudformation stack resources. Which of the following can be used to help in the debugging process?
Choose 2 answers from the options below

A. UseCloudtrail to debugall the API call's sent by the Cloudformation stack.
B. Usethe AWS CloudFormation console to view the status of yourstack.
C. Seethe logs in the/var/log directory for Linux instances
D. UseAWSConfig to debug all the API call's sent by the Cloudformation stack.

Answer: BC

Explanation:
The AWS Documentation mentions
Use the AWS Cloud Formation console to view the status of your stack. In the console, you can view a list of stack events while your stack is being created, updated, or
deleted. From this list, find the failure event and then view the status reason for that event.
For Amazon CC2 issues, view the cloud-init and cfn logs. These logs are published on the Amazon CC2 instance in the /var/log/ directory. These logs capture processes and command outputs while AWS Cloud Formation is setting up your instance. For Windows, view the L~C2Configure service and cfn logs in %ProgramFiles%\Amazon\CC2ConfigService and C:\cfn\log.
For more information on Cloudformation Troubleshooting, please visit the below URL:
http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/troubleshooting.html

NEW QUESTION 198
As part of your continuous deployment process, your application undergoes an I/O load performance test before it is deployed to production using new AMIs. The application uses one Amazon EBS PIOPS volume per instance and requires consistent I/O performance.
Which of the following must be carried out to ensure that I/O load performance tests yield the correct results in a repeatable manner?

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

Answer: B

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

NEW QUESTION 199
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 202
You are managing an application that contains Go as the front end, MongoDB for document management and is hosted on a relevant Web server. You pre-bake AMI'S with the latest version of the Web server, then user the User Data section to setup the application. You now have a change to the underlying Operating system version and need to deploy that accordingly. How can this be done
in the easiest way possible.

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

Answer: C

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

NEW QUESTION 205
You are using Elastic Beanstalk for your development team. You are responsible for deploying multiple versions of your application. How can you ensure, in an ideal way, that you don't cross the application version limit in Elastic beanstalk?

A. Createa lambda function to delete the older versions.
B. Createa script to delete the older versions.
C. UseAWSConfig to delete the older versions
D. Uselifecyle policies in Elastic beanstalk

Answer: D

Explanation:
The AWS Documentation mentions
Each time you upload a new version of your application with the Clastic Beanstalk console or the CB CLI, Elastic Beanstalk creates an application version. If you don't delete versions that you no longer use, you will eventually reach the application version limit and be unable to create new versions of that application.
You can avoid hitting the limit by applying an application version lifecycle policy to your applications.
A lifecycle policy tells Clastic Beanstalk to delete application versions that are old, or to delete application versions when the total number of versions for an application exceeds a specified number.
For more information on Clastic Beanstalk lifecycle policies please see the below link:
• http://docs.aws.a mazon.com/elasticbeanstalk/latest/dg/appl ications-lifecycle.html

NEW QUESTION 210
When deploying applications to Elastic Beanstalk, which of the following statements is false with regards to application deployment

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

Answer: B

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

NEW QUESTION 211
You are planning on configuring logs for your Elastic Load balancer. At what intervals does the logs get produced by the Elastic Load balancer service. Choose 2 answers from the options given below

A. 5minutes
B. 60minutes
C. 1 minute
D. 30seconds

Answer: AB

Explanation:
The AWS Documentation mentions
Clastic Load Balancing publishes a log file for each load balancer node at the interval you specify. You can specify a publishing interval of either 5 minutes or 60 minutes when you enable the access log for your load balancer. By default. Elastic Load Balancing publishes logs at a 60-minute interval.
For more information on Elastic load balancer logs please see the below link: http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/access-log-collection.html

NEW QUESTION 212
Which of the following tools is available to send logdatafrom EC2 Instances.

A. CloudWatch LogsAgent
B. CloudWatchAgent
C. Logsconsole.
D. LogsStream

Answer: A

Explanation:
The AWS Documentation mentions the following
The CloudWatch Logs agent provides an automated way to send log data to Cloud Watch Logs from Amazon L~C2 instances. The agent is comprised of the following components:
A plug-in to the AWS CLI that pushes log data to CloudWatch Logs.
A script (daemon) that initiates the process to push data to CloudWatch Logs.
Acron job that ensures that the daemon is always running. For more information on Cloudwatch logs Agent, please see the below link:
http://docs.aws.a mazon.com/AmazonCloudWatch/latest/logs/AgentRefe re nee. htm I

NEW QUESTION 215
Which of the following is not a supported platform for the Elastic beanstalk service

A. Java
B. AngularJS
C. PHP
D. .Net

Answer: B

Explanation:
Below are the supported platforms for Elastic beanstalk
Platforms
- Packer Builder
- Single Container Docker
- Multicontainer Docker
- Preconfigured Docker
- Go
- Java SE
- Java with Tomcat
- .NET on Windows Server with IIS
- Node.js
- PHP
- Python
- Ruby

For more information on Elastic beanstalk, please visit the below URL:
http://docs.aws.a mazon.com/elasticbeanstalk/latest/dg/concepts.platforms. htm I

NEW QUESTION 218
By default in Opswork, how many application versions can you rollback up to?

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

Answer: D

Explanation:
The AWS Documentation mentions the following Restores the previously deployed app version. For example, if you have deployed the app three times and then run Rollback, the server will serve the app from the second deployment. If you run Rollback again, the server will serve the app from the first deployment. By default, AWS OpsWorks Stacks stores the five most recent deployments, which allows you to roll back up to four versions. If you exceed the number of stored

versions, the command fails and leaves the oldest version in place.
For more information on Opswork app deployment, please visit the below U RL: http://docs.aws.amazon.com/opsworks/latest/userguide/workingapps-deploying.html

NEW QUESTION 221
You have an AWS OpsWorks Stack running Chef Version 11.10. Your company hosts its own proprietary cookbook on Amazon S3, and this is specified as a custom cookbook in the stack. You want to use an open-source cookbook located in an external Git repository. What tasks should you perform to enable the use of both custom cookbooks?

A. Inthe AWS OpsWorks stack settings, enable Berkshel
B. Create a new cookbook with aBerksfile that specifies the other two cookbook
C. Configure the stack to usethis new cookbook.
D. Inthe OpsWorks stack settings add the open source project's cookbook details inaddition to your cookbook.
E. Contactthe open source project's maintainers and request that they pull your cookbookinto their
F. Update the stack to use their cookbook.
G. Inyour cookbook create an S3 symlink object that points to the open sourceproject's cookbook.

Answer: A

Explanation:
To use an external cookbook on an instance, you need a way to install it and manage any dependencies. The preferred approach is to implement a cookbook that supports a dependency manager named Berkshelf. Berkshelf works on Amazon CC2 instances, including AWS OpsWorks Stacks instances, but it is also designed to work with Test Kitchen and Vagrant.
For more information on Opswork and Berkshelf, please visit the link:
• http://docs.aws.a mazon.com/opsworks/latest/userguide/cookbooks-101 -opsworks- berkshelf.htm l

NEW QUESTION 223
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::Instanceresource 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 224
You are a Devops Engineer and are designing an Opswork stack in AWS. The company has some custom recipes that are part of their on-premise Chef configuration. These same recipes need to be run whenever an instance is launched in Opsworks. Which of the following steps need to be carried out to to ensure this requirement gets fulfilled. Choose 2 answers from the options given below

A. Ensurethe custom cookbooks option is set in Opswork stack.
B. Ensurethe custom cookbooks option is set in Opswork layer.
C. Ensurethe recipe is placed as part of the Setup Lifecycle event as part of the Layersetting.
D. Ensurethe recipe is placed as part of the Setup Lifecycle event as part of the Stacksetting.

Answer: AC

Explanation:
The AWS Documentation mentions the below
Each layer has a set of built-in recipes assigned to each lifecycle event, although some layers lack Undeploy recipes. When a lifecycle event occurs on an instance,
AWS OpsWorks Stacks runs the appropriate set of recipes for the associated layer.
The below diagram shows a snapshot in the stack settings where you can enable custom cookbooks

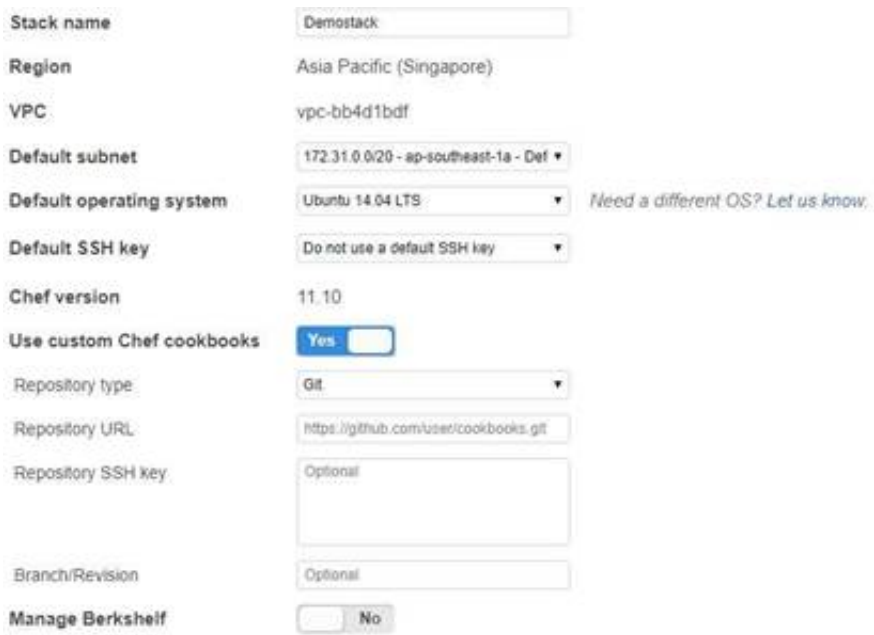

For more information on automating recipe's, please visit the below URL: http://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-assigningcustom.html

NEW QUESTION 228
You are a Devops Engineer for your company. The company has a number of Cloudformation templates in AWS. There is a concern from the IT Security department and they want to know who all use the Cloudformation stacks in the company's AWS account. Which of the following can be done to take care of this security concern?

A. EnableCloudwatch events for each cloudformation stack to track the resource creationevents.
B. EnableCloudtrail logs so that the API calls can be recorded
C. EnableCloudwatch logs for each cloudformation stack to track the resource creationevents.
D. ConnectSQS and Cloudformation so that a message is published for each resource createdin the Cloudformation stack.

Answer: B

Explanation:
This is given as a best practice in the AWS documentation
AWS CloudTrail tracks anyone making AWS Cloud Formation API calls in your AWS account. API calls are logged whenever anyone uses the AWS Cloud Formation API, the AWS Cloud Formation console, a back-end console, or AWS CloudFormation AWS CLI commands.
Enable logging and specify an Amazon S3 bucket to store the logs. That way, if you ever need to, you can audit who made what AWS CloudFormation call in your account
For more information on the best practises, please visit the below URL: http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/best-practices.html

NEW QUESTION 233
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

DOP-C01 Practice Exam Features:

* DOP-C01 Questions and Answers Updated Frequently

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

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

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