

Amazon-Web-Services

Exam Questions DOP-C01

AWS Certified DevOps Engineer- Professional



NEW QUESTION 1

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 using an AWS CloudFormation template
- 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 template
- D. Spin up a second AWS CloudFormation stack
- 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 Scaling
- G. Setup your Auto Scaling group policies to scale based on the number of requests per second as well as the current customer load time
- H. •>/D- Re-deploy your application using an Auto Scaling template
- 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 EC2 instances available to handle the load for your application. You create collections of EC2 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 your group 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 your group 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/WhatIsAutoScaling.html>

NEW QUESTION 2

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

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

Answer: BD

Explanation:

Option B makes sense because maybe the max servers is low hence the application cannot handle the peak load.

Option D helps in ensuring Autoscaling can scale the group on the right metrics.

For more information on Autoscaling health checks, please refer to the below document link: from AWS <http://docs.aws.amazon.com/autoscaling/latest/userguide/healthcheck.html>

NEW QUESTION 3

Management has reported an increase in the monthly bill from Amazon Web Services, and they are extremely concerned with this increased cost. Management has asked you to determine the exact cause of this increase. After reviewing the billing report, you notice an increase in the data transfer cost. How can you provide management with a better insight into data transfer use?

- A. Update your Amazon CloudWatch metrics to use five-second granularity, which will give better detailed metrics that can be combined with your billing data to pinpoint anomalies.
- B. Use Amazon CloudWatch Logs to run a map-reduce on your logs to determine high usage and data transfer.
- C. Deliver custom metrics to Amazon CloudWatch per application that breaks down application data transfer into multiple, more specific data points.
- D- Using Amazon CloudWatch metrics, pull your Elastic Load Balancing outbound data transfer metrics monthly, and include them with your billing report to show which application is causing higher bandwidth usage.

Answer: C

Explanation:

You can publish your own metrics to CloudWatch using the AWS CLI or an API. You can view statistical graphs of your published metrics with the AWS Management Console.

CloudWatch stores data about a metric as a series of data points. Each data point has an associated time stamp. You can even publish an aggregated set of data points called a statistic set.

If you have custom metrics specific to your application, you can give a breakdown to the management on the exact issue.

Option A won't be sufficient to provide better insights.

Option B is an overhead when you can make the application publish custom metrics. Option D is invalid because just the ELB metrics will not give the entire picture.

For more information on custom metrics, please refer to the below document link: from AWS

<http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.html>

NEW QUESTION 4

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

Answer: D

Explanation:

IAM 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 IAM Roles please refer to the below link:
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

NEW QUESTION 5

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

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

Answer: BC

Explanation:

The AWS::AutoScaling::AutoScalingGroup resource supports an UpdatePolicy attribute. This is used to define how an Auto Scaling group resource is updated when an update to the CloudFormation 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 it is not efficient to use CloudFormation to use Elastic Beanstalk. Option D is invalid because this is an inefficient process to tear down stacks when there are stack policies available. For more information on AutoScaling Rolling Updates please refer to the below link:
• <https://aws.amazon.com/premiumsupport/knowledge-center/auto-scaling-group-rolling-updates/>

NEW QUESTION 6

You have a complex system that involves networking, IAM 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-control.
- 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 for finer granularity 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::Stack resource in your template to reference other templates. For more information on CloudFormation best practices please refer to the below link: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html>

NEW QUESTION 7

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

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

Answer: BDE

Explanation:

You can use Cloud Watch Logs to monitor applications and systems using log data

Cloud Watch Logs uses your log data for monitoring; so, no code changes are required. For example, you can monitor application logs for specific literal terms (such as "NullPointerException") or count the number of occurrences of a literal term at a particular position in log data (such as "404" status codes in an Apache access log). When the term you are searching for is found. Cloud Watch Logs reports the data to a CloudWatch metric that you specify. Log data is encrypted while in transit and while it is at rest

For more information on Cloudwatch logs please refer to the below link: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html>

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

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

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

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

NEW QUESTION 8

You have enabled Elastic Load Balancing HTTP health checking. After looking at the AWS Management Console, you see that all instances are passing health checks, but your customers are reporting that your site is not responding. What is the cause?

- A. The HTTP health checking system is misreporting due to latency in inter-instance metadata synchronization.
- B. The health check in place is not sufficiently evaluating the application function.
- C. The application is returning a positive health check too quickly for the AWS Management Console to respond.
- D. Latency in DNS resolution is interfering with Amazon EC2 metadata retrieval.

Answer: B

Explanation:

You need to have a custom health check which will evaluate the application functionality. Its not enough using the normal health checks. If the application functionality does not work and if you don't have custom health checks, the instances will still be deemed as healthy.

If you have custom health checks, you can send the information from your health checks to Auto Scaling so that Auto Scaling can use this information. For example, if you determine that an instance is not functioning as expected, you can set the health status of the instance to Unhealthy. The next time that Auto Scaling performs a health check on the instance, it will determine that the instance is unhealthy and then launch a replacement instance

For more information on Autoscaling health checks, please refer to the below document link: from AWS

<http://docs.aws.amazon.com/autoscaling/latest/userguide/healthcheck.html>

NEW QUESTION 9

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

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

Answer: A

Explanation:

Please check the following AWS DOCs which provides details on the scenario. Check the example of "configure".

? <https://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-events.html> You can use the Configure Lifecycle event

This event occurs on all of the stack's instances when one of the following occurs:

- An instance enters or leaves the online state.
- You associate an Elastic IP address with an instance or disassociate one from an instance.
- You attach an Elastic Load Balancing load balancer to a layer, or detach one from a layer. Ensure the Opswork layer uses a custom Cookbook

2. Toggle **Use custom Chef cookbooks** to **Yes**.



The screenshot shows the AWS OpsWorks console configuration for a stack. The 'Use custom Chef cookbooks' toggle is set to 'Yes'. Below this, the 'Repository type' is set to 'Git', the 'Repository URL' is 'https://github.com/aws-labs/op...', and the 'Repository SSH key' is 'Optional'. The 'Branch/Revision' is also set to 'Optional'. Under 'Stack color', there are several color swatches. At the bottom, there are two diagrams of an 'Amazon RDS Multi-AZ' instance within an 'AWS Region'.

For more information on Opswork stacks, please refer to the below document link: from AWS

- http://docs.aws.amazon.com/opsworks/latest/userguide/welcome_classic.html

NEW QUESTION 10

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

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

Answer: CD

Explanation:

Option A is invalid because there is no "aws git.push" command

Option B is invalid because there is no AWS::EC2::Container resource type.

Elastic Beanstalk supports the deployment of web applications from Docker containers. With Docker containers, you can define your own runtime environment. You can choose your own platform, programming language, and any application dependencies (such as package managers or tools), that aren't supported by other platforms. Docker containers are self-contained and include all the configuration information and software your web application requires to run.

For more information on Docker and Elastic beanstalk please refer to the below link:

• http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker.html

When you launch an instance in Amazon EC2, you have the option of passing user data to the instance that can be used to perform common automated configuration tasks and even run scripts after the instance starts. You can pass two types of user data to Amazon EC2: shell scripts and cloud-init directives. You can also pass this data into the launch wizard as plain text, as a file (this is useful for launching instances using the command line tools), or as base64-encoded text (for API calls). For more information on EC2 User data please refer to the below link:

• <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/user-data.html>

Note: "git aws.push" with CB CLI 2.x - see a forum thread at <https://forums.aws.amazon.com/thread.jspa?threadID=583202&messageID=582979>. Basically, this is a predecessor to the newer "eb deploy" command in CB CLI 3.1. This question kept in order to be consistent with exam.

NEW QUESTION 10

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

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

Answer: A

Explanation:

Elastic Beanstalk supports the deployment of web applications from Docker containers. With Docker containers, you can define your own runtime environment. You can choose your own platform, programming language, and any application dependencies (such as package managers or tools), that aren't supported by other platforms. Docker containers are self-contained and include all the configuration information and software your web application requires to run.

By using Docker with Elastic Beanstalk, you have an infrastructure that automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring.

For more information on Docker and Elastic beanstalk please refer to the below link:

• http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker.html

NEW QUESTION 13

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

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

Answer: A

Explanation:

An instance profile is a container for an IAM role that you can use to pass role information to an EC2 instance when the instance starts.

Option B is invalid because you can associate a role with an instance

Option C and D are invalid because using users or user groups is not a pre-requisite For more information on instance profiles, please visit the link:

• http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2-instance-profiles.html

NEW QUESTION 15

You have an application running on Amazon EC2 in an Auto Scaling group. Instances are being bootstrapped dynamically, and the bootstrapping takes over 15 minutes to complete. You find that instances are reported by Auto Scaling as being In Service before bootstrapping has completed. You are receiving application alarms related to new instances before they have completed bootstrapping, which is causing confusion. You find the cause: your application monitoring tool is polling the Auto Scaling Service API for instances that are In Service, and creating alarms for new previously unknown instances. Which of the following will ensure

that new instances are not added to your application monitoring tool before bootstrapping is completed?

- A. Create an Auto Scaling group lifecycle hook to hold the instance in a pending: wait state until your bootstrapping is complete
- B. Once bootstrapping is complete, notify Auto Scaling to complete the lifecycle hook and move the instance into a pending:proceed state.
- C. Use the default Amazon Cloud Watch application metrics to monitor your application's health
- D. Configure an Amazon SNS topic to send these Cloud Watch alarms to the correct recipients.
- E. Tag all instances on launch to identify that they are in a pending state
- F. Change your application monitoring tool to look for this tag before adding new instances, and then use the Amazon API to set the instance state to 'pending' until bootstrapping is complete.
- G. Increase the desired number of instances in your Auto Scaling group configuration to reduce the time it takes to bootstrap future instances.

Answer: A

Explanation:

Auto Scaling lifecycle hooks enable you to perform custom actions as Auto Scaling launches or terminates instances. For example, you could install or configure software on newly launched instances, or download log files from an instance before it terminates. After you add lifecycle hooks to your Auto Scaling group, they work as follows:

1. Auto Scaling responds to scale out events by launching instances and scale in events by terminating instances.
2. Auto Scaling puts the instance into a wait state (Pending:Wait or Terminating:Wait). The instance remains in this state until either you tell Auto Scaling to continue or the timeout period ends.

For more information on rolling updates, please visit the below link:

- <http://docs.aws.amazon.com/autoscaling/latest/userguide/lifecycle-hooks.htm>

NEW QUESTION 19

You've been tasked with improving the current deployment process by making it easier to deploy and reducing the time it takes. You have been tasked with creating a continuous integration (CI) pipeline that can build AMI's. Which of the below is the best manner to get this done. Assume that at max your development team will be deploying builds 5 times a week.

- A. Use a dedicated EC2 instance with an EBS Volume
- B. Download and configure the code and then create an AMI out of that.
- C. Use OpsWorks to launch an EBS-backed instance, then use a recipe to bootstrap the instance, and then have the CI system use the CreateImage API call to make an AMI from it.
- D. Upload the code and dependencies to Amazon S3, launch an instance, download the package from Amazon S3, then create the AMI with the CreateSnapshot API call
- E. Have the CI system launch a new instance, then bootstrap the code and dependencies on that instance, and create an AMI using the CreateImage API call.

Answer: D

Explanation:

Since the number of builds is just a few times a week, there are many open source systems such as Jenkins which can be used as CI based systems.

Jenkins can be used as an extensible automation server, Jenkins can be used as a simple CI server or turned into the continuous delivery hub for any project.

For more information on the Jenkins CI tool please refer to the below link:

- <https://jenkins.io/>

Option A and C are partially correct, but since you just have 5 deployments per week, having separate instances which consume costs is not required. Option B is partially correct, but again having a separate system such as Opsworks for such a low number of deployments is not required.

NEW QUESTION 22

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

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

Answer: AB

Explanation:

Amazon Cloud Watch monitors your Amazon Web Services (AWS) resources and the applications you run on AWS in real time. You can use Cloud Watch to collect and track metrics, which are variables you can measure for your resources and applications. Cloud Watch alarms send notifications or automatically make changes to the resources you are monitoring based on rules that you define.

For more information on AWS Cloudwatch, please refer to the below document link: from AWS

- <http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/WhatIsCloudWatch.htm> | Amazon Cloud Watch uses Amazon SNS to send email. First, create and subscribe to an SNS topic.

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

For more information on AWS Cloudwatch and SNS, please refer to the below document link: from AWS

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

NEW QUESTION 24

What is web identity federation?

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

Answer: B

Explanation:

With web identity federation, you don't need to create custom sign-in code or manage your own user identities. Instead, users of your app can sign in using a well-

known identity provider (IdP) — such as Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible IdP, receive an authentication token, and then exchange that token for temporary security credentials in AWS that map to an IAM role with permissions to use the resources in your AWS account. Using an IdP helps you keep your AWS account secure, because you don't have to embed and distribute long-term security credentials with your application. For more information on Web Identity federation please refer to the below link:
http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html

NEW QUESTION 27

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

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

Answer: D

Explanation:

Amazon Elasticsearch Service makes it easy to deploy, operate, and scale Elasticsearch for log analytics, full text search, application monitoring, and more. Amazon

Elasticsearch Service is a fully managed service that delivers Elasticsearch's easy-to-use APIs and real-time capabilities along with the availability, scalability, and security required by production workloads. The service offers built-in integrations with Kibana, Logstash, and AWS services including Amazon Kinesis Firehose, AWS Lambda, and Amazon CloudWatch so that you can go from raw data to actionable insights quickly. For more information on Elastic Search, please refer to the below link:

- <https://aws.amazon.com/elasticsearch-service/>

NEW QUESTION 32

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

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

Answer: C

Explanation:

Option A is invalid, because CloudWatch will not store the logs indefinitely and secondly it won't be the cost-effective option.

Option B is invalid, because it won't serve the purpose of regularly retrieving the most recent logs for troubleshooting. You will need to pay more to retrieve the logs faster from this storage.

Option D is invalid, because it is not an ideal or cost-effective option.

You can define lifecycle configuration rules for objects that have a well-defined lifecycle. For example: if you are uploading periodic logs to your bucket, your application might need these logs for a week or a month after creation, and after that you might want to delete them.

Some documents are frequently accessed for a limited period of time. After that, these documents are less frequently accessed. Over time, you might not need real-time access to these objects, but your organization or regulations might require you to archive them for a longer period and then optionally delete them later.

You might also upload some types of data to Amazon S3 primarily for archival purposes, for example digital media archives, financial and healthcare records, raw genomics sequence data, long-term database backups, and data that must be retained for regulatory compliance.

For more information on Lifecycle management please refer to the below link: <http://docs.aws.amazon.com/AmazonS3/latest/dev/object-lifecycle-mgmt.html>

Note:

Option C is the cheapest option, but CloudWatch can store logs indefinitely or between 10 years and one day.

"Log Retention—By default, logs are kept indefinitely and never expire. You can adjust the retention policy for each log group, keeping the indefinite retention, or choosing a retention period between 10 years and one day." <https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html>

NEW QUESTION 33

You have deployed a CloudFormation template which is used to spin up resources in your account. Which of the following status in CloudFormation represents a failure.

- A. UPDATE_COMPLETE_CLEANUP_IN_PROGRESS
- B. DELETE_COMPLETE
- C. ROLLBACK_IN_PROGRESS
- D. UPDATE_IN_PROGRESS

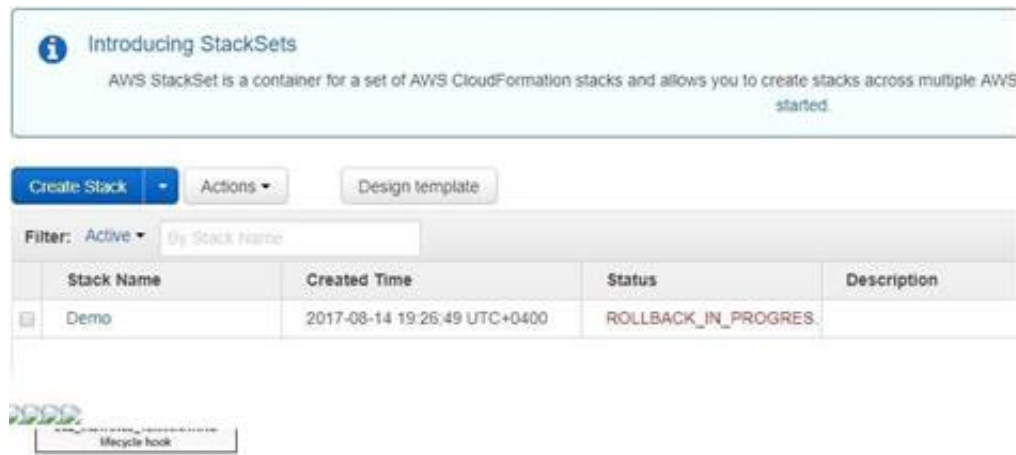
Answer: C

Explanation:

AWS CloudFormation provisions and configures resources by making calls to the AWS services that are described in your template.

After all the resources have been created, AWS CloudFormation reports that your stack has been created. You can then start using the resources in your stack. If stack creation fails, AWS CloudFormation rolls back your changes by deleting the resources that it created.

The below snapshot from CloudFormation shows what happens when there is an error in the stack creation.



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

NEW QUESTION 35

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

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

Answer: B

Explanation:

Since the time required to spin up an instance is required to be fast, its better to create an AMI rather than use User Data. When you use User Data, the script will be

run during boot up, and hence this will be slower.

An Amazon Machine Image (AMI) provides the information required to launch an instance, which is a virtual server in the cloud. You specify an AMI when you launch

an instance, and you can launch as many instances from the AMI as you need. You can also launch instances from as many different AMIs as you need.

For more information on the AMI, please refer to the below link:

- <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIs.html>

NEW QUESTION 39

Your company uses AWS to host its resources. They have the following requirements

- 1) Record all API calls and Transitions
- 2) Help in understanding what resources are there in the account
- 3) Facility to allow auditing credentials and logins

Which services would suffice the above requirements

- A. AWS Config, CloudTrail, IAM Credential Reports
- B. CloudTrail, IAM Credential Reports, AWS Config
- C. CloudTrail, AWS Config, IAM Credential Reports
- D. AWS Config, IAM Credential Reports, CloudTrail

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.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html>

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. With Config, you can review changes in configurations and relationships between AWS resources, dive into detailed resource configuration histories, and determine your overall compliance against the configurations specified in your internal guidelines. This enables you to simplify compliance auditing, security analysis, change management, and operational troubleshooting. For more information on the config service, please visit the below URL:

- <https://aws.amazon.com/config/>

You can generate and download a credential report that lists all users in your account and the status of their various credentials, including passwords, access keys, and MFA devices. You can get a credential report from the AWS Management Console, the AWS SDKs and Command Line Tools, or the IAM API. For more information on Credentials Report, please visit the below URL:

- http://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_getting-report.html

NEW QUESTION 44

You need the absolute highest possible network performance for a cluster computing application. You already selected homogeneous instance types supporting 10 gigabit enhanced networking, made sure that your workload was network bound, and put the instances in a placement group. What is the last optimization you can make?

- A. Use 9001 MTU instead of 1500 for Jumbo Frames, to raise packet body to packet overhead ratios.
- B. Segregate the instances into different peered VPCs while keeping them all in a placement group, so each one has its own Internet Gateway.
- C. Bake an AMI for the instances and relaunch, so the instances are fresh in the placement group and do not have noisy neighbors.
- D. Turn off SYN/ACK on your TCP stack or begin using UDP for higher throughput.

Answer: A

Explanation:

Jumbo frames allow more than 1500 bytes of data by increasing the payload size per packet, and thus increasing the percentage of the packet that is not packet overhead. Fewer packets are needed to send the same amount of usable data. However, outside of a given AWS region (CC2-Classic), a single VPC, or a VPC peering connection, you will experience a maximum path of 1500 MTU. VPN connections and traffic sent over an Internet gateway are limited to 1500 MTU. If packets are over 1500 bytes, they are fragmented, or they are dropped if the Don't Fragment flag is set in the IP header. For more information on Jumbo Frames, please visit the below URL:
http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/network_mtu.htm#jumbo_frame_instances

NEW QUESTION 47

Your application's Auto Scaling Group scales up too quickly, too much, and stays scaled when traffic decreases. What should you do to fix this?

- A. Set a longer cooldown period on the Group, so the system stops overshooting the target capacity
- B. The issue is that the scaling system doesn't allow enough time for new instances to begin servicing requests before measuring aggregate load again.
- C. Calculate the bottleneck or constraint on the compute layer, then select that as the new metric, and set the metric thresholds to the bounding values that begin to affect response latency.
- D. Raise the CloudWatch Alarms threshold associated with your autoscaling group, so the scaling takes more of an increase in demand before beginning.
- E. Use larger instances instead of lots of smaller ones, so the Group stops scaling out so much and wasting resources as the OS level, since the OS uses a higher proportion of resources on smaller instances.

Answer: B

Explanation:

The ideal case is that the right metric is not being used for the scale up and down.

Option A is not valid because it mentions that the cooldown is not happening when the traffic decreases, that means the metric threshold for the scale down is not occurring in Cloudwatch

Option C is not valid because increasing the Cloudwatch alarm metric will not ensure that the instances scale down when the traffic decreases.

Option D is not valid because the question does not mention any constraints that points to the instance size. For an example on using custom metrics for scaling in and out, please follow the below link for a use case.

- <https://blog.powerupcloud.com/aws-autoscaling-based-on-database-query-custom-metrics-f396c16e5e6a>

NEW QUESTION 52

Your CTO has asked you to make sure that you know what all users of your AWS account are doing to change resources at all times. She wants a report of who is doing what over time, reported to her once per week, for as broad a resource type group as possible. How should you do this?

- A. Create a global AWS CloudTrail Trail
- B. Configure a script to aggregate the log data delivered to S3 once per week and deliver this to the CTO.
- C. Use CloudWatch Events Rules with an SNS topic subscribed to all AWS API call
- D. Subscribe the CTO to an email type delivery on this SNS Topic.
- E. Use AWS IAM credential reports to deliver a CSV of all uses of IAM UserTokens overtime to the CTO.
- F. Use AWS Config with an SNS subscription on a Lambda, and insert these changes over time into a DynamoDB table
- G. Generate reports based on the contents of this table.

Answer: A

Explanation:

AWS CloudTrail is an AWS service that helps you enable governance, compliance, and operational and risk auditing of your AWS account. Actions taken by a user, role, or an AWS service are recorded as events in CloudTrail. Events include actions taken in the AWS Management Console, AWS Command Line Interface, and AWS SDKs and APIs.

Visibility into your AWS account activity is a key aspect of security and operational best practices. You can use CloudTrail to view, search, download, archive, analyze, and respond to account activity across your AWS infrastructure. You can identify who or what took which action, what resources were acted upon, when the event occurred, and other details to help you analyze and respond to activity in your AWS account.

For more information on Cloudtrail, please visit the below URL:

- <http://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html>

NEW QUESTION 54

You meet once per month with your operations team to review the past month's data. During the meeting, you realize that 3 weeks ago, your monitoring system which pings over HTTP from outside AWS recorded a large spike in latency on your 3-tier web service API. You use DynamoDB for the database layer, ELB, EBS, and EC2 for the business logic tier, and SQS, ELB, and EC2 for the presentation layer. Which of the following techniques will NOT help you figure out what happened?

- A. Check your CloudTrail log history around the spike's time for any API calls that caused slowness.
- B. Review CloudWatch Metrics for one minute interval graphs to determine which components slowed the system down.
- C. Review your ELB access logs in S3 to see if any ELBs in your system saw the latency.
- D. Analyze your logs to detect bursts in traffic at that time.

Answer: B

Explanation:

The Cloudwatch metric retention is as follows. If the data points are of a one minute interval, then the graphs will not be available in Cloudwatch

- Data points with a period of less than 60 seconds are available for 3 hours. These data points are high-resolution custom metrics.
- Data points with a period of 60 seconds (1 minute) are available for 15 days
- Data points with a period of 300 seconds (5 minute) are available for 63 days
- Data points with a period of 3600 seconds (1 hour) are available for 455 days (15 months) For more information on Cloudwatch metrics, please visit the below URL:
• http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/cloudwatch_concepts.html

NEW QUESTION 57

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 60

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

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

Answer: D

Explanation:

Since the issue is occuring after the new feature has been added, it could be relevant to the new feature.

Enabling Cloudtrail will just monitor all the API calls of all services and will not benefit the cause.

The monitoring of CPU utilization will just reverify that there is an issue but will not help pinpoint the issue.

The Elastic Load Balancer logs will also just reverify that there is an issue but will not help pinpoint the issue.

For more information on custom Cloudwatch metrics, please refer to the below link:

<http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.html>

NEW QUESTION 65

Which of the following is the default deployment mechanism used by Elastic Beanstalk when the application is created via Console or EBCLI?

- A. All at Once
- B. Rolling Deployments
- C. Rolling with additional batch
- D. Immutable

Answer: B

Explanation:

The AWS documentation mentions

AWS Elastic Beanstalk provides several options for how deployments are processed, including deployment policies (All at once. Rolling, Rolling with additional batch,

and Immutable) and options that let you configure batch size and health check behavior during deployments. By default, your environment uses rolling deployments

if you created it with the console or EB CLI, or all at once deployments if you created it with a different client (API, SDK or AWS CLI).

For more information on Elastic Beanstalk deployments, please refer to the below link:

- <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.rolling-version-deploy.html>

NEW QUESTION 67

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 72

You have launched a cloudformation template, but are receiving a failure notification after the template was launched. What is the default behavior of Cloudformation in such a case

- A. It will rollback all the resources that were created up to the failure point.

- B. It will keep all the resources that were created up to the failure point.
- C. It will prompt the user on whether to keep or terminate the already created resources
- D. It will continue with the creation of the next resource in the stack

Answer: A

Explanation:

The AWS Documentation mentions

AWS Cloud Formation ensures all stack resources are created or deleted as appropriate. Because AWS CloudFormation treats the stack resources as a single unit,

they must all be created or deleted successfully for the stack to be created or deleted. If a resource cannot be created, AWS CloudFormation rolls the stack back and automatically deletes any resources that were created.

For more information on Cloudformation, please refer to the below link: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacks.html>

NEW QUESTION 73

When building a multicontainer Docker platform using Elastic Beanstalk, which of the following is required

- A. DockerFile to create custom images during deployment
- B. Prebuilt Images stored in a public or private online image repository.
- C. Kubernetes to manage the docker containers.
- D. RedHatOpensift to manage the docker containers.

Answer: B

Explanation:

This is a special note given in the AWS Documentation for Multicontainer Docker platform for Elastic Beanstalk

Building custom images during deployment with a Dockerfile is not supported by the multicontainer Docker platform on Elastic Beanstalk. Build your images and deploy them to an online repository before creating an Elastic Beanstalk environment.

For more information on Multicontainer Docker platform for Elastic Beanstalk, please refer to the below link:

http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_ecs.html

NEW QUESTION 78

Which of the following credentials types are supported by AWSCodeCommit? Select 3 Options

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

Answer: ABD

Explanation:

The AWS documentation mentions

I AM supports AWS CodeCommit with three types of credentials:

Git credentials, an IAM -generated user name and password pair you can use to communicate with AWS CodeCommit repositories over HTTPS.

SSH keys, a locally generated public-private key pair that you can associate with your IAM user to communicate with AWS CodeCommit repositories over SSH.

AWS access keys, which you can use with the credential helper included with the AWS CLI to communicate with AWS CodeCommit repositories over HTTPS.

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_ssh-keys.html

NEW QUESTION 81

Which of the following is the right sequence of initial steps in the deployment of application revisions using Code Deploy

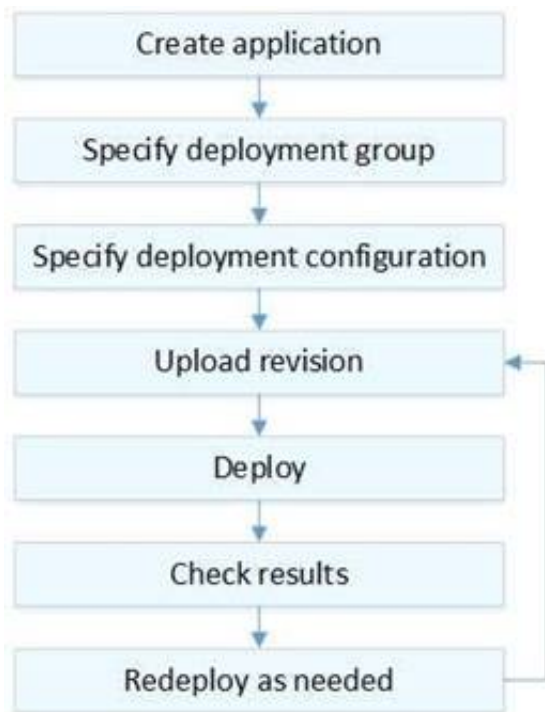
- 1) Specify deployment configuration
- 2) Upload revision
- 3) Create application
- 4) Specify deployment group

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

Answer: C

Explanation:

The below diagram from the AWS documentation shows the deployment steps



For more information on the deployment steps please refer to the below link:

- <http://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-steps.html>

NEW QUESTION 85

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

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

Answer: B

Explanation:

The AWS Documentation mentions

Auto Scaling enables you to put an instance that is in the InService state into the Standby state, update or troubleshoot the instance, and then return the instance to service. Instances that are on standby are still part of the Auto Scaling group, but they do not actively handle application traffic.

For more information on the standby state please refer to the below link:

- <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-enter-exit-standby.html>

NEW QUESTION 86

You have an Opswork stack setup in AWS. You want to install some updates to the Linux instances in the stack. Which of the following can be used to publish those updates. Choose 2 answers from the options given below

- A. Create and start new instances to replace your current online instance
- B. Then delete the current instances.
- C. Use Auto-scaling to launch new instances and then delete the older instances
- D. On Linux-based instances in Chef 11.10 or older stacks, run the Update Dependencies stack command
- E. Delete the stack and create a new stack with the instances and their relevant updates

Answer: AC

Explanation:

As per AWS documentation.

By default, AWS OpsWorks Stacks automatically installs the latest updates during setup, after an instance finishes booting. AWS OpsWorks Stacks does not automatically install updates after an instance is online, to avoid interruptions such as restarting application servers. Instead, you manage updates to your online instances yourself, so you can minimize any disruptions.

We recommend that you use one of the following to update your online instances.

- Create and start new instances to replace your current online instances. Then delete the current instances.

The new instances will have the latest set of security patches installed during setup.

- On Linux-based instances in Chef 11.10 or older stacks, run the Update Dependencies stack command, which installs the current set of security patches and other updates on the specified instances.

More information is available at: <https://docs.aws.amazon.com/opsworks/latest/userguide/workingsecurity-updates.html>

NEW QUESTION 91

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 Elastic Beanstalk-specific JSON file that describes how to deploy a set of Docker containers as an Elastic 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 95

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

Configuration Updates

The following settings control how changes to the environment's instances are propagated.

Rolling update type:	<div> <div>Rolling based on Health</div> <div>Disabled</div> <div>Rolling based on Health</div> <div>Rolling based on Time</div> <div>Immutable</div> </div>	Learn more
Maximum batch size:	Number of instances that should be modified at any given time.	
Minimum instances in service:	1	The minimum number of instances that should be in service at any given time.
Pause time:	<div> <div>0</div> Hour <div>0</div> Minutes <div>0</div> Seconds </div> The time to wait between changes to a batch of environments. Must be an hour or less.	

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 bootstrap 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 99

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

- A. Create a powershell script which uses the AWS CLI to get the volumes and then run the script as a cron job.
- B. Use the AWSConfig service to create a snapshot of the AWS Volumes
- C. Use the AWS CodeDeploy service to create a snapshot of the AWS Volumes
- D. Use Cloudwatch Events to trigger the snapshots of EBS Volumes

Answer: D

Explanation:

The best is to use the inbuilt service from Cloudwatch, as Cloud watch Events to automate the creation of CBS Snapshots. With Option A, you would be restricted to running the powrshell script on Windows machines and maintaining the script itself And then you have the overhead of having a separate instance just to run that script.

When you go to Cloudwatch events, you can use the Target as EC2 CreateSnapshot API call as shown below.

Create rules to invoke Targets based on Events happening in your AWS environment.

Event Source

Build or customize an Event Pattern or set a Schedule to invoke Targets.

Event Pattern ☒ Schedule ☒

Fixed rate of: 5 Minutes

Cron expression: */5 * * * *

Learn more about CloudWatch Events schedules

* Show sample event(s)

* Required

Targets

Select Target to invoke when an event matches your Event Pattern or when schedule is triggered.

Lambda function

- CodePipeline
- EC2 CreateSnapshot API call
- EC2 RebootInstances API call
- EC2 StopInstances API call
- EC2 TerminateInstances API call
- EC2 Task
- Event bus in another AWS account
- Kinesis stream

Cancel Configure targets

The AWS Documentation mentions

Amazon Cloud Watch Cvents delivers a near real-time stream of system events that describe changes in Amazon Web Services (AWS) resources. Using simple rules

that you can quickly set up, you can match events and route them to one or more target functions or streams. Cloud Watch Cvents becomes aware of operational changes as they occur. Cloud Watch Cvents responds to these operational changes and takes corrective action as necessary, by sending messages to respond to the environment, activating functions, making changes, and capturing state information. For more information on Cloud watch Cvents, please visit the below U RL:

- <http://docs.aws.amazon.com/AmazonCloudWatch/latest/events/WhatIsCloudWatchEvents.html>

NEW QUESTION 104

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

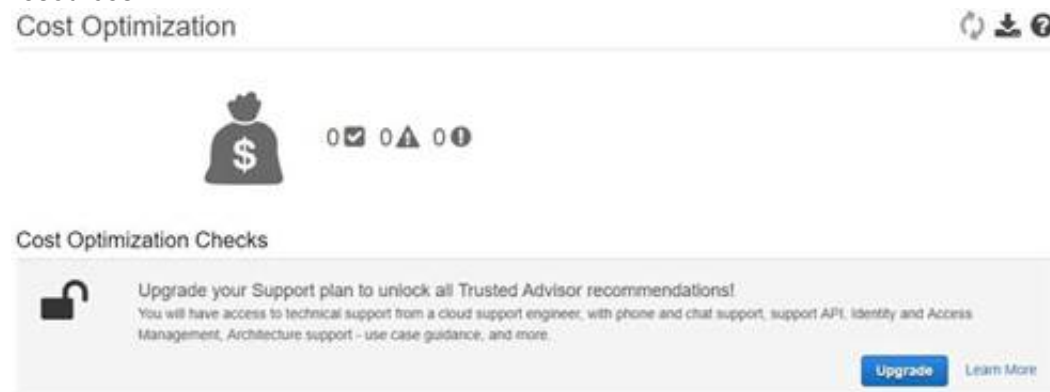
Explanation:

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/AlarmThatSendsEmail.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 URL:

- <https://aws.amazon.com/premiumsupport/trustedadvisor/>

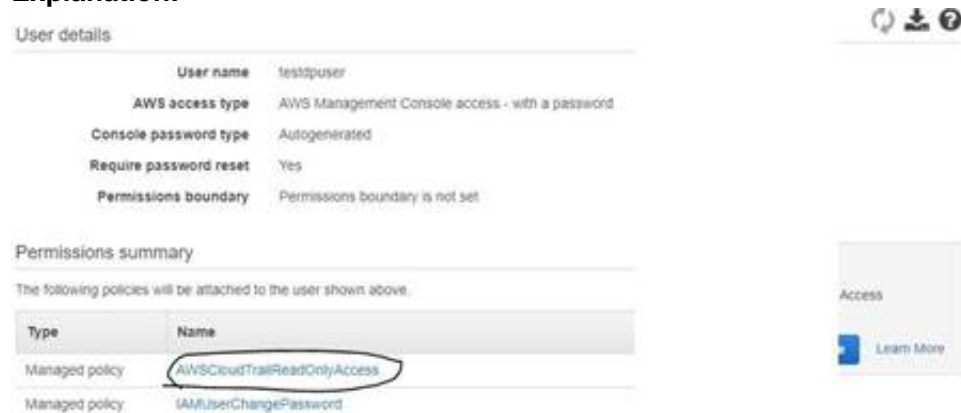
NEW QUESTION 105

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

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

Answer: B

Explanation:



The AWS Documentation clearly mentions the below

AWS CloudTrail is an AWS service that helps you enable governance, compliance, and operational and risk auditing of your AWS account. Actions taken by a user, role, or an AWS service are recorded as events in CloudTrail. Events include actions taken in the AWS Management Console, AWS Command Line Interface, and AWS SDKs and APIs.

For more information on Cloudtrail, please visit the below URL:

- <http://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html>

NEW QUESTION 108

Your company has the requirement to set up instances running as part of an Autoscaling Group. Part of the requirement is to use Lifecycle hooks to setup custom based software's and do the necessary configuration on the instances. The time required for this setup might take an hour, or might finish before the hour is up. How should you setup lifecycle hooks for the Autoscaling Group. Choose 2 ideal actions you would include as part of the lifecycle hook.

- A. Configure the lifecycle hook to record heartbeat
- B. If the hour is up, restart the timeout period.
- C. Configure the lifecycle hook to record heartbeat
- D. If the hour is up, choose to terminate the current instance and start a new one
- E. Ifthe software installation and configuration is complete, then restart the time period.
- F. If the software installation and configuration is complete, then send a signal to complete the launch of the instance.

Answer: AD

Explanation:

The AWS Documentation provides the following information on lifecycle hooks

By default, the instance remains in a wait state for one hour, and then Auto Scaling continues the launch or terminate process (Pending: Proceed or Terminating: Proceed). If you need more time, you can restart the timeout period by recording a heartbeat. If you finish before the timeout period ends, you can complete the

lifecycle action, which continues the launch or termination process

For more information on AWS Lifecycle hooks, please visit the below URL:

- <http://docs.aws.amazon.com/autoscaling/latest/userguide/lifecycle-hooks.html>

NEW QUESTION 112

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

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

Answer: C

Explanation:

Generally you should refrain from giving high level permissions and give only the required permissions. In this case option C fits well by just providing the relevant access which is required.

For more information on IAM please see the below link:

- <https://aws.amazon.com/iam/>

NEW QUESTION 117

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

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

Answer: A

Explanation:

The AWS Documentation mentions

The optional Mappings section matches a key to a corresponding set of named values. For example, if you want to set values based on a region, you can create a mapping that uses the region name as a key and contains the values you want to specify for each specific region. You use the `Fn::FindInMap` intrinsic function to retrieve values in a map.

For more information on mappings please refer to the below link:

- <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/mappings-section-structure.html>

NEW QUESTION 120

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 125

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

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

the LDAP credentials the application can use the 1AM temporary credentials to access the appropriate AWS service.


Answer: BC

Explanation:

When you have the need for an in-premise environment to work with a cloud environment, you would normally have 2 artefacts for authentication purposes

- An identity store - So this is the on-premise store such as Active Directory which stores all the information for the user's and the groups they belong to.
- An identity broker - This is used as an intermediate agent between the on-premise location and the cloud environment. In Windows you have a system known as Active Directory Federation services to provide this facility.

Hence in the above case, you need to have an identity broker which can work with the identity store and the Security Token service in AWS. An example diagram of how this works from the AWS documentation is given below.



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

NEW QUESTION 128

A user is accessing RDS from an application. The user has enabled the Multi AZ feature with the MS SQL RDS DB. During a planned outage how will AWS ensure that a switch from DB to a standby replica will not affect access to the application?

- A. RDS will have an internal IP which will redirect all requests to the new DB
- B. RDS uses DNS to switch over to stand by replica for seamless transition
- C. The switch over changes Hardware so RDS does not need to worry about access
- D. RDS will have both the DBs running independently and the user has to manually switch over

Answer: B

Explanation:

Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable.

In case of an infrastructure failure (for example, instance hardware failure, storage failure, or network disruption), Amazon RDS performs an automatic failover to the standby, so that you can resume database operations as soon as the failover is complete.

And as per the AWS documentation, the canonical name is changed to the standby DB when the primary one fails.

Q: What happens during Multi-AZ failover and how long does it take?

"Failover is automatically handled by Amazon RDS so that you can resume database operations as quickly as possible without administrative intervention. When failing over, Amazon RDS simply flips the canonical name record (CNAMR) for your DB instance to point at the standby, which is in turn promoted to become the new primary. We encourage you to follow best practices and implement database connection retry at the application layer".

<https://aws.amazon.com/rds/faqs/>

Based on this, RDS Multi-AZ will use DNS to create the CNAMR and hence B is the right option. For more information on RDS Multi-AZ please visit the link:

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html>

NEW QUESTION 131

You are currently using SQS 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. Use Kinesis as a buffer stream for message bodies
- 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. */
- E. Use SQS's support for message partitioning and multi-part uploads on Amazon S3.
- F. Use AWS EFS as a shared pool storage medium
- G. Store filesystem pointers to the files on 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 134

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 stack creation will wait until one service is created and launched
- B. The user can use the `HoldCondition` resource to wait for the creation of the other dependent resources
- C. The user can use the `DependentCondition` resource to hold the creation of the other dependent resources
- D. The user can use the `WaitCondition` resource 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.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-properties-waitcondition.html>

NEW QUESTION 139

A gaming company adopted AWS Cloud Formation to automate load-testing of their games. 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) PostgreSQL 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 the load-testing AWS Cloud Formation stack and define a dependency relation between master and replica via the `Depends On` attribute.
- B. Define an update policy to prevent deletion of the Amazon RDS database after the AWS Cloud Formation stack is deleted.
- C. Define a deletion policy of type `Retain` for the Amazon RDS resource to assure that the RDS database is not deleted with the AWS Cloud Formation stack.
- D. Define a deletion policy of type `Snapshot` for the Amazon RDS resource to assure that the RDS database can be restored after the AWS Cloud Formation stack is deleted.
- E. Define automated backups with a backup retention period of 30 days for the Amazon RDS database and perform point-in-time recovery of the database after the AWS CloudFormation stack is deleted.

Answer: CD

Explanation:

With the `DeletionPolicy` 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 EC2 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-attribute-deletionpolicy.html>

NEW QUESTION 144

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

- A. Create separate dashboards in every region
- B. Register instances running on different regions to CloudWatch
- C. Have one single dashboard to report metrics to CloudWatch from different region
- D. This is not possible

Answer: C

Explanation:

You can monitor AWS resources in multiple regions using a single Cloud Watch dashboard. For example, you can create a dashboard that shows CPU utilization for an

EC2 instance located in the us-west-2 region with your billing metrics, which are located in the us-east-1 region.

For more information on Cloudwatch dashboard, please refer to the below url

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

NEW QUESTION 146

A web-startup runs its very successful social news application on Amazon EC2 with an Elastic Load Balancer, an Auto-Scaling group of Java/Tomcat application servers, and DynamoDB as data store. The main web application best runs on m2.xlarge instances since it is highly memory-bound. Each new deployment requires semi-automated creation and testing of a new AMI for the application servers which takes quite a while and is therefore only done once per week. Recently, a new chat feature has been implemented in Node.js and waits to be integrated in the architecture. First tests show that the new component is CPU bound because the company has some experience with using Chef, they decided to streamline the deployment process and use AWS Ops Works as an application life cycle tool to simplify management of the application and reduce the deployment cycles. What configuration in AWS Ops Works is necessary to integrate the new chat module in the most cost-efficient and flexible way?

- A. Create one AWS Ops Works stack, create one AWS Ops Works layer, create one custom recipe
- B. Create one AWS Ops Works stack create two AWS Ops Works layers create one custom recipe
- C. Create two AWS Ops Works stacks create two AWS Ops Works layers create one custom recipe
- D. Create two AWS Ops Works stacks create two AWS Ops Works layers create two custom recipes

Answer: B

Explanation:

You can just have one Opswork stack and multiple layers, one for Node js and the other for the standard application.

An AWS OpsWorks Stack defines the configuration of your entire application: the load balancers, server software, database, etc. You control every part of the stack by building layers that define the software packages deployed to your instances and other configuration details such as Elastic IPs and security groups. You can also deploy your software onto layers by identifying the repository and optionally using Chef Recipes to automate everything Chef can do, such as creating directories and users, configuring databases, etc. You can use OpsWorks Stacks' built-in automation to scale your application and automatically recover from instance failures.

You can control who can view and manage the resources that are used by your application, including ssh access to the instances that your application uses.

For more information on Ops work, please visit the below URL

- <https://aws.amazon.com/opsworks/stacks/faqs/>

NEW QUESTION 149

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. Store all the data in S3 so that it can be more cost effective
- B. Store your data in Amazon S3, and use lifecycle policies to archive to Amazon Glacier
- C. Store your data in an EBS volume, and use lifecycle policies to archive to Amazon Glacier.
- D. Store your data in Amazon S3, and use lifecycle policies to archive to S3-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 STANDARD_IA (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.amazon.com/AmazonS3/latest/dev/object-lifecycle-mgmt.html>

NEW QUESTION 151

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

- A. Create dynamic scaling policies to scale up on Monday and scale down on Friday
- B. Create a scheduled policy to scale up on Friday and scale down on Monday
- C. Create a scheduled policy to scale up on Monday and scale down on Friday
- D. Manually add instances to the Autoscaling Group on Monday and remove them on Friday

Answer: C

Explanation:

The AWS Documentation mentions the following for Scheduled scaling

Scaling based on a schedule allows you to scale your application in response to predictable load changes. For example, every week the traffic to your web application starts to increase on Wednesday, remains high on Thursday, and starts to decrease on Friday. You can plan your scaling activities based on the predictable traffic patterns of your web application.

For more information on scheduled scaling for Autoscaling, please visit the below URL

- http://docs.aws.amazon.com/autoscaling/latest/userguide/schedule_time.html

NEW QUESTION 154

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

Choose 2 answers from the options below.

- A. Create Docker containers for the custom application components.
- B. Use OpsWorks to deploy the Docker containers.
- C. Use Elastic Beanstalk to deploy the Docker containers.
- D. Use CloudFormation to deploy the Docker containers.

Answer: AC

Explanation:

The AWS documentation states the following for Docker containers on Elastic Beanstalk

Elastic Beanstalk supports the deployment of web applications from Docker containers. With Docker containers, you can define your own runtime environment. You can choose your own platform, programming language, and any application dependencies (such as package managers or tools), that aren't supported by other platforms. Docker containers are self-contained and include all the configuration information and software your web application requires to run.

For more information on Docker containers and Elastic Beanstalk, please visit the below URL

http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker.html

NEW QUESTION 156

Your company has recently extended its datacenter into a VPC on AWS. There is a requirement for on-premise users to manage AWS resources from the AWS console. You don't want to create IAM users for them again. Which of the below options will fit your needs for authentication?

- A. Use OAuth 2.0 to retrieve temporary AWS security credentials to enable your members to sign in to the AWS Management Console.
- B. Use Web Identity Federation to retrieve AWS temporary security credentials to enable your members to sign in to the AWS Management Console.
- C. Use your on-premise SAML 2 O-compliant identity provider (IDP) to grant members federated access to the AWS Management Console via the AWS

singlesign-on (SSO) endpoint.

D. Use your on-premises SAML2.0-compliant identity provider (IDP) to retrieve temporary security credentials to enable members to sign in to the AWS Management Console.

Answer: C

Explanation:

You can use a role to configure your SAML 2.0-compliant IdP and AWS to permit your federated users to access the AWS Management Console. The role grants the user permissions to carry out tasks in the console.

For more information on AWS SAML, please visit the below URL

- http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_enable-console-saml.html

NEW QUESTION 160

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 165

You are a DevOps Engineer in your company. You have been instructed to ensure there is an automated backup solution in place for EBS Volumes. These snapshots need to be retained only for a period of 20 days. How can you achieve this requirement in an efficient manner?

- A. Use the aws ec2 create-volume API to create a snapshot of the EBS Volume
- B. Use the describe-volumes API to see those snapshots which are greater than 20 days and then delete them accordingly using the delete-volume API call.
- C. Use Lifecycle policies to push the EBS Volumes to Amazon Glacie
- D. Then use further lifecycle policies to delete the snapshots after 20 days.
- E. Use Lifecycle policies to push the EBS Volumes to Amazon S3. Then use further lifecycle policies to delete the snapshots after 20 days.
- F. Use Amazon Data Lifecycle Manager to automate the process.

Answer: D

Explanation:

Use Amazon Data Lifecycle Manager (Amazon DLM) to automate the creation, retention, and deletion of snapshots taken to back up your Amazon EBS volumes. Automating snapshot management helps you to:

- Protect valuable data by enforcing a regular backup schedule. Retain backups as required by auditors or internal compliance.
- Reduce storage costs by deleting outdated backups.

For more information, please check the below AWS Docs:

- <https://docs.aws.amazon.com/AWSC2/latest/UserGuide/snapshot-lifecycle.html>

NEW QUESTION 166

You are designing a CloudFormation stack which involves the creation of a web server and a database server. You need to ensure that the web server in the stack gets created after the database server is created. How can you achieve this?

- A. Ensure that the database server is defined first and before the web server in the CloudFormation template
- B. The stack creation normally goes in order to create the resources.
- C. Ensure that the database server is defined as a child of the web server in the CloudFormation template.
- D. Ensure that the web server is defined as a child of the database server in the CloudFormation template.
- E. Use the DependsOn attribute to ensure that the database server is created before the web server.

Answer: D

Explanation:

The AWS Documentation mentions

With the DependsOn attribute you can specify that the creation of a specific resource follows another. When you add a DependsOn attribute to a resource, that resource is created only after the creation of the resource specified in the DependsOn attribute.

For more information on the DependsOn attribute, please visit the below URL <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-dependson.html>

NEW QUESTION 168

Which of the following is false when it comes to using the Elastic Load balancer with OpsWorks stacks?

- A. You can attach only one load balancer to a layer.
- B. A Classic Load Balancer can span across AWS OpsWorks Stacks layers.
- C. Each load balancer can handle only one layer.
- D. You need to create the load balancer before hand and then attach it to the OpsWorks stack.

Answer: B

Explanation:

The AWS Documentation mentions the following

To use Elastic Load Balancing with a stack, you must first create one or more load balancers in the same region by using the Elastic Load Balancing console, CLI, or API. You should be aware of the following:

You can attach only one load balancer to a layer. Each 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 Elastic Load Balancer with Opswork,

please visit the below url <http://docs.aws.amazon.com/opsworks/latest/userguide/layers-elb.html>

NEW QUESTION 173

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

- A. Insert a timestamp for each record and then scan the entire table for the timestamp as per the last 2 hours.
- B. Create another DynamoDB table with the records modified in the last 2 hours.
- C. Use DynamoDB streams to monitor the changes in the DynamoDB table.
- D. Transfer the records to S3 which were modified in the last 2 hours

Answer: C

Explanation:

The AWS Documentation mentions the following

A DynamoDB stream is an ordered flow of information about changes to items in an Amazon DynamoDB table. When you enable a stream on a table, DynamoDB captures information about every modification to data items in the table.

Whenever an application creates, updates, or deletes items in the table, DynamoDB Streams writes a stream record with the primary key attribute(s) of the items that were modified. A stream record contains information about a data modification to a single item in a DynamoDB table. You can configure the stream so that the stream records capture additional information, such as the "before" and "after" images of modified items.

For more information on DynamoDB streams, please visit the below URL: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.html>

NEW QUESTION 176

Which of the following are the basic stages of a CI/CD Pipeline. Choose 3 answers from the options below

- A. SourceControl
- B. Build
- C. Run
- D. Production

Answer: ABD

Explanation:

The below diagram shows the stages of a typical CI/CD pipeline

Continuous Integration



For more information on AWS Continuous Integration, please visit the below URL: <https://d.awsstatic.com/whitepapers/DevOps/practicing-continuous-integration-continuous-delivery-on-AWS.pdf>

NEW QUESTION 177

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. Use CloudTrail to debug all the API call's sent by the CloudFormation stack.
- B. Use the AWS CloudFormation console to view the status of your stack.
- C. See the logs in the /var/log directory for Linux instances
- D. Use AWSConfig 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 EC2 issues, view the cloud-init and cfn logs. These logs are published on the Amazon EC2 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/AWSCloudFormation/latest/UserGuide/troubleshooting.html>

NEW QUESTION 178

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

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

Answer: C

Explanation:

Answer - C

This is given in the AWS Documentation list-stack-resources

Description

Returns descriptions of all resources of the specified stack.

For deleted stacks, ListStackResources returns resource information for up to 90 days after the stack has been deleted.

See also: AWS API Documentation

See 'aws help' for descriptions of global parameters.

list-stack-resources is a paginated operation. Multiple API calls may be issued in order to retrieve the entire data set of results. You can disable pagination by providing the `—no-paginate` argument. When using `—output text` and the `—query` argument on a paginated response, the `—query` argument must extract data from the results of the following query expressions: StackResourceSummaries For more information on the CLI command, please visit the below URL:

<http://docs.aws.amazon.com/cli/latest/reference/cloudformation/list-stack-resources.html>

NEW QUESTION 182

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

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

Answer: C

Explanation:

When you add an Availability Zone to your load balancer. Clastic Load Balancing creates a load balancer node in the Availability Zone. Load balancer nodes accept traffic from clients and forward requests to the healthy registered instances in one or more Availability Zones.

For more information on adding AZ's to CLB, please refer to the below U RL:

<http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/enable-disable-az.html>

NEW QUESTION 184

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

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

Answer: C

Explanation:

The AWS Documentation mentions the below

Amazon Kinesis makes it easy to collect, process, and analyze real-time, streaming data so you can get timely insights and react quickly to new information. Amazon

Kinesis offers key capabilities to cost effectively process streaming data at any scale, along with the flexibility to choose the tools that best suit the requirements of your application. With Amazon Kinesis, you can ingest real-time data such as application logs, website clickstreams, IoT telemetry data, and more into your databases, data lakes and data warehouses, or build your own real-time applications using this data.

Amazon Kinesis enables you to process and analyze data as it arrives and respond in real-time instead of having to wait until all your data is collected before the processing can begin.

For more information on AWS Kinesis please see the below link:

- <https://aws.amazon.com/kinesis/>

NEW QUESTION 188

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

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

Answer: AC

Explanation:

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

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

For more information on monitoring DynamoDB please see the below link:

- <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/monitoring-cloudwatch.html>

NEW QUESTION 193

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

- A. The application can be bundled in a zip file
- B. Can include parent directories
- C. Should not exceed 512 MB in size
- D. Can be a war file which can be deployed to the application server

Answer: B

Explanation:

The AWS Documentation mentions

When you use the AWS Elastic 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 Elastic Beanstalk please see the below link: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/applications-sourcebundle.html>

NEW QUESTION 196

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

- A. Create a new ELB and attach the AutoScaling Group to the ELB
- B. Create a new launch configuration with the updated AMI and associate it with the AutoScaling group
- C. Increase the size of the group to six and when instances become healthy revert to three.
- D. Manually terminate the instances with the older launch configuration.
- E. Update the launch configuration instead of updating the AutoScaling Group

Answer: B

Explanation:

An Auto Scaling group is associated with one launch configuration at a time, and you can't modify a launch configuration after you've created it. To change the launch configuration for an Auto Scaling group, you can use an existing launch configuration as the basis for a new launch configuration and then update the Auto Scaling group to use the new launch configuration.

After you change the launch configuration for an Auto Scaling group, any new instances are launched using the new configuration options, but existing instances are not affected.

Then to ensure the new instances are launched, change the size of the AutoScaling Group to 6 and once the new instances are launched, change it back to 3.

For more information on instances scale-in process and Auto Scaling Group's termination policies please view the following link:

- <https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-instance-termination.html#default-termination-policy> For more information on changing the launch configuration please see the below link:
- <http://docs.aws.amazon.com/autoscaling/latest/userguide/change-launch-config.html>

NEW QUESTION 197

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

- A. Create a new CloudTrail trail with one new S3 bucket to store the logs and with the global services option selected
- B. Use IAM roles, S3 bucket policies and MultiFactor Authentication (MFA) Delete on the S3 bucket that stores your log
- C. V
- D. Create a new CloudTrail with one new S3 bucket to store the log
- E. Configure SNS to send log file delivery notifications to your management system
- F. Use IAM roles and S3 bucket policies on the S3 bucket that stores your logs.
- G. Create a new CloudTrail trail with an existing S3 bucket to store the logs and with the global services option selected
- H. Use S3 ACLs and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs.
- I. Create three new CloudTrail trails with three new S3 buckets to store the logs one for the AWS Management console, one for AWS SDKs and one for command line tools. Use IAM roles and S3 bucket policies on the S3 buckets that store your logs.

Answer: A

Explanation:

AWS Identity and Access Management (IAM) is integrated with AWS CloudTrail, a service that logs AWS events made by or on behalf of your AWS account.

CloudTrail logs authenticated AWS API calls and also AWS sign-in events, and collects this event information in files that are delivered to Amazon S3 buckets.

You need to ensure that all services are included. Hence option B is partially correct.

Option B and D is wrong because it just adds an overhead for having 3 S3 buckets and SNS notifications.

For more information on CloudTrail, please visit the below URL:

- <http://docs.aws.amazon.com/IAM/latest/UserGuide/cloudtrail-integration.html>

NEW QUESTION 199

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.amazon.com/elasticbeanstalk/latest/dg/concepts.platforms.html>

NEW QUESTION 204

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 URL: <http://docs.aws.amazon.com/opsworks/latest/userguide/workingapps-deploying.html>

NEW QUESTION 205

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

- A. Add another stack with the database layer and attach it to the application stack.
- B. Configure SSL so that the instance can communicate with the database
- C. Add the appropriate driver packages to ensure the application can work with the database
- D. Configure database tags for the Opswork application layer

Answer: C

Explanation:

The AWS documentation mentions the below point Important

For Linux stacks, if you want to associate an Amazon RDS service layer with your app, you must add the appropriate driver package to the associated app server layer, as follows:

1. Click Layers in the navigation pane and open the app server's Recipes tab.
2. Click Edit and add the appropriate driver package to OS Packages. For example, you should specify mysql if the layer contains Amazon Linux instances and mysql-client if the layer contains Ubuntu instances.
3. Save the changes and redeploy the app.

For more information on Opswork app connectivity, please visit the below URL: <http://docs.aws.amazon.com/opsworks/latest/userguide/workingapps-connectdb.html>

NEW QUESTION 209

Your application requires long-term storage for backups and other data that you need to keep readily available but with lower cost. Which S3 storage option should

you use?

- A. AmazonS3 Standard- Infrequent Access
- B. S3Standard
- C. Glacier
- D. ReducedRedundancy Storage

Answer: A

Explanation:

The AWS Documentation mentions the following

Amazon S3 Standard - Infrequent Access (Standard - IA) is an Amazon S3 storage class for data that is accessed less frequently, but requires rapid access when needed. Standard - IA offers the high durability, throughput, and low latency of Amazon S3 Standard, with a low per GB storage price and per GB retrieval fee.

For more information on S3 Storage classes, please visit the below URL:

- <https://aws.amazon.com/s3/storage-classes/>

NEW QUESTION 211

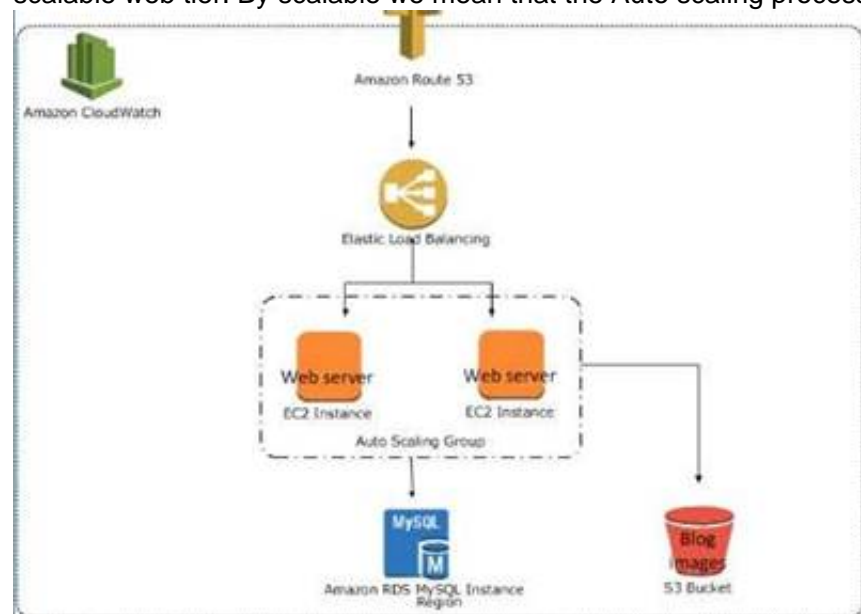
A company is building a two-tier web application to serve dynamic transaction-based content. The data tier is leveraging an Online Transactional Processing (OLTP) database. What services should you leverage to enable an elastic and scalable web tier?

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

Answer: A

Explanation:

The question mentioned a scalable web tier and not a database tier. So Option C, D and B are already automated eliminated, since we do not need a database option. The below example shows an Elastic Load balancer connected to 2 EC2 instances connected via Auto Scaling. This is an example of an elastic and scalable web tier. By scalable we mean that the Auto scaling process will increase or decrease the number of EC2 instances as required.



For more information on best practices for AWS Cloud applications, please visit the below URL:

https://dO.awsstatic.com/whitepapers/AWS_Cloud_Best_Practices.pdf

NEW QUESTION 214

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

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

Answer: C

Explanation:

The AWS CloudHSM service helps you meet corporate, contractual and regulatory compliance requirements for data security by using dedicated Hardware Security

Module (HSM) appliances within the AWS cloud. With CloudHSM, you control the encryption keys and cryptographic operations performed by the HSM.

Option D is wrong with the CloudHSM option because of the ephemeral volume which this is temporary storage

For more information on cloudhsm, please refer to the link:

- <https://aws.amazon.com/cloudhsm/>

NEW QUESTION 215

You are working with a customer who is using Chef Configuration management in their data center. Which service is designed to let the customer leverage existing

Chef recipes in AWS?

- A. AmazonSimple Workflow Service
- B. AWSElastic Beanstalk
- C. AWSCloudFormation
- D. AWSOpsWorks

Answer: D

Explanation:

AWS OpsWorks is a configuration management service that helps you configure and operate applications of all shapes and sizes using Chef. You can define the application's architecture and the specification of each component including package installation, software configuration and resources such as storage. Start from templates for common technologies like application servers and databases or build your own to perform any task that can be scripted. AWS OpsWorks includes automation to scale your application based on time or load and dynamic configuration to orchestrate changes as your environment scales.

For more information on Opswork, please visit the link:

- <https://aws.amazon.com/opsworks/>

NEW QUESTION 218

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 aBerkshelf file that specifies the other two cookbook
- C. Configure the stack to use this new cookbook.
- D. Inthe OpsWorks stack settings add the open source project's cookbook details in addition to your cookbook.
- E. Contactthe open source project's maintainers and request that they pull your cookbook into 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.amazon.com/opsworks/latest/userguide/cookbooks-101-opsworks-berkshelf.html>

NEW QUESTION 222

A group of developers in your organization want to migrate their existing application into Elastic Beanstalk and want to use Elastic load Balancing and Amazon SQS. They are currently using a custom application server. How would you deploy their system to Elastic Beanstalk?

- A. Configurean Elastic Beanstalk platform using AWS OpsWorks deploy it to Elastic Beanstalkand run a script that creates a load balancer and an Amazon SQS queue.
- B. Usea Docker container that has the third party application server installed on itand that creates the load balancer and an Amazon SQS queue using the applicationsource bundle feature.
- C. Createa custom Elastic Beanstalk platform that contains the third party applicationserver and runs a script that creates a load balancer and an Amazon SQS queue.
- D. Configurean AWS OpsWorks stack that installs the third party application server andcreates a load balancer and an Amazon SQS queue and then deploys it to ElasticBeanstalk.

Answer: B

Explanation:

Below is the documentation on Elastic beanstalk and Docker

Elastic Beanstalk supports the deployment of web applications from Docker containers. With Docker containers, you can define your own runtime environment. You

can choose your own platform, programming language, and any application dependencies (such as package managers or tools), that aren't supported by other platforms. Docker containers are self-contained and include all the configuration information and software your web application requires to run.

For more information on Elastic beanstalk and Docker, please visit the link: http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker.html

NEW QUESTION 225

What is the amount of time that Opswork stacks services waits for a response from an underlying instance before deeming it as a failed instance?

- A. 1minute.
- B. 5minutes.
- C. 20minutes.
- D. 60minutes

Answer: B

Explanation:

The AWS Documentation mentions

Every instance has an AWS OpsWorks Stacks agent that communicates regularly with the service. AWS OpsWorks Stacks uses that communication to monitor instance health. If an agent does not communicate with the service for more than approximately five minutes, AWS OpsWorks Stacks considers the instance to have failed.

For more information on the Auto healing feature, please visit the below URL: <http://docs.aws.amazon.com/opsworks/latest/userguide/workinginstances-auto>

healing.html

NEW QUESTION 228

Which of the following file needs to be included along with your source code binaries when your application uses the EC2/On-Premises compute platform, and deploy it using the AWS Code Deploy service.

- A. appspec.yml
- B. appconfig.yml
- C. appspecjson
- D. appconfigjson

Answer: A

Explanation:

The AWS Documentation mentions the below

The application specification file (AppSpec file) is a YAML-formatted file used by AWS CodeDeploy to determine:

what it should install onto your instances from your application revision in Amazon S3 or GitHub. which lifecycle event hooks to run in response to deployment lifecycle events. An AppSpec file must be named appspec.yml and it must be placed in the root of an application's source code's directory structure. Otherwise, deployments will fail. For more information on the appspec file, please visit the below URL:

<http://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file.html>

Note: If you deploying your code on AWS Lambda compute platform. An AppSpec file can be YAML- formatted or JSON-formatted. You can also enter the contents of an AppSpec file directly into AWS CodeDeploy console when you create a deployment.

However, this question is about along with your source code binaries on an EC2/On-Premises Compute Platform. So, an AppSpec file must be a YAML-formatted file named appspec.yml and it must be placed in the root of the directory structure of an application's source code. Otherwise, deployments fail.

NEW QUESTION 230

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 creation events.
- B. EnableCloudtrail logs so that the API calls can be recorded
- C. EnableCloudwatch logs for each cloudformation stack to track the resource creation events.
- D. ConnectSQS and Cloudformation so that a message is published for each resource created in 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/AWSCloudFormation/latest/UserGuide/best-practices.html>

NEW QUESTION 234

Which of the following are ways to secure data at rest and in transit in AWS. Choose 3 answers from the options given below

- A. Encrypt all EBS volumes attached to EC2 Instances
- B. Use server side encryption for S3
- C. Use SSL/HTTPS when using the Elastic Load Balancer
- D. Use IOPS volumes when working with EBS volumes on EC2 Instances

Answer: ABC

Explanation:

The AWS documentation mentions the following

Amazon CBS encryption offers you a simple encryption solution for your EBS 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 Data protection refers to protecting data while in-transit (as it travels to and from Amazon S3) and at rest (while it is stored on disks in Amazon S3 data centers). You can protect data in transit by using SSL or by using client-side encryption. You have the following options of protecting data at rest in Amazon S3.

Use Server-Side encryption - You request Amazon S3 to encrypt your object before saving it on disks in its data centers and decrypt it when you download the objects.

Use Client-Side Encryption - You can encrypt data client-side and upload the encrypted data to Amazon S3. In this case, you manage the encryption process, the encryption keys, and related tools. You can create a load balancer that uses the SSL/TLS protocol for encrypted connections (also known as SSL offload). This feature enables traffic encryption between your load balancer and the clients that initiate HTTPS sessions, and for connections between your load balancer and your EC2 instances. For more information on securing data at rest, please refer to the below link:

• <https://d03wsstatic.com/whitepapers/aws-securing-data-at-rest-with-encryption.pdf>

NEW QUESTION 238

Your company is planning on using the available services in AWS to completely automate their integration, build and deployment process. They are planning on using AWS CodeBuild to build their artefacts. When using CodeBuild, which of the following files specifies a collection of build commands that can be used by the service during the build process.

- A. appspec.yml
- B. buildspec.yml

C. buildspec.xml
D. appspec.json

Answer: B

Explanation:

The AWS documentation mentions the following

AWS CodeBuild currently supports building from the following source code repository providers. The source code must contain a build specification (build spec) file,

or the build spec must be declared as part of a build project definition. A buildspec is a collection of build commands and related settings, in YAML format, that AWS

CodeBuild uses to run a build.

For more information on AWS CodeBuild, please refer to the below link: <http://docs.aws.amazon.com/codebuild/latest/userguide/planning.html>

NEW QUESTION 240

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

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

100% Actual & Verified — Instant Download, Please Click
[Order The DOP-C01 Practice Test Here](#)