

Amazon-Web-Services

Exam Questions DOP-C01

AWS Certified DevOps Engineer- Professional



NEW QUESTION 1

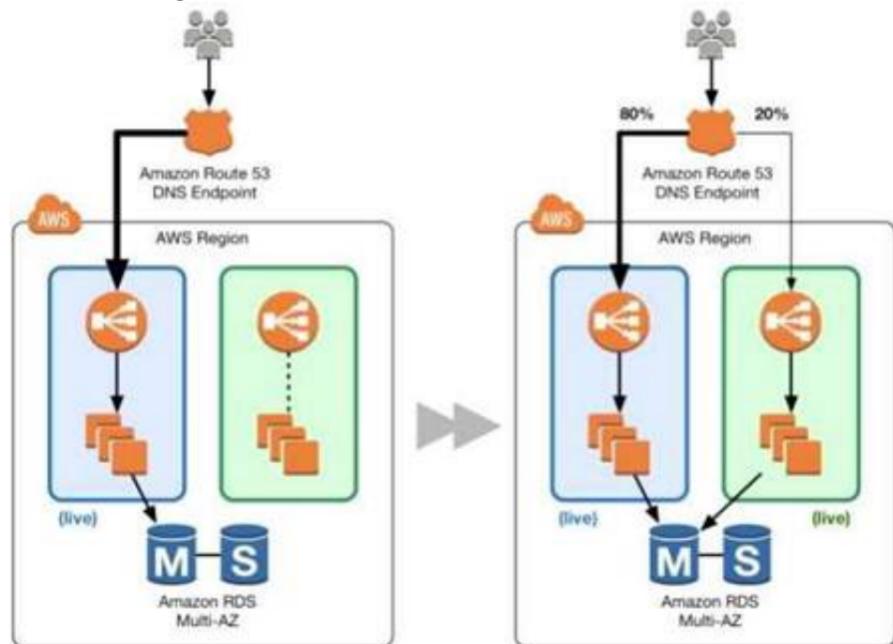
Your application is currently running on Amazon EC2 instances behind a load balancer. Your management has decided to use a Blue/Green deployment strategy. How should you implement this for each deployment?

- A. Set up Amazon Route 53 health checks to fail over from any Amazon EC2 instance that is currently being deployed to.
- B. Using AWS CloudFormation, create a test stack for validating the code, and then deploy the code to each production Amazon EC2 instance.
- C. Create a new load balancer with new Amazon EC2 instances, carry out the deployment, and then switch DNS over to the new load balancer using Amazon Route 53 after testing.
- D. Launch more Amazon EC2 instances to ensure high availability, de-register each Amazon EC2 instance from the load balancer, upgrade it, and test it, and then register it again with the load balancer.

Answer: C

Explanation:

The below diagram shows how this can be done



- 1) First create a new ELB which will be used to point to the new production changes.
 - 2) Use the Weighted Route policy for Route53 to distribute the traffic to the 2 ELB's based on a 80- 20% traffic scenario. This is the normal case, the % can be changed based on the requirement.
 - 3) Finally when all changes have been tested, Route53 can be set to 100% for the new ELB.
- Option A is incorrect because this is a failover scenario and cannot be used for Blue green deployments. In Blue Green deployments, you need to have 2 environments running side by side. Option B is incorrect, because you need to have a production stack with the changes which will run side by side. Option D is incorrect because this is not a blue green deployment scenario. You cannot control which users will go to the new EC2 instances. For more information on blue green deployments, please refer to the below document link: from AWS https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 2

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

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

Answer: C

Explanation:

The ideal way is to create a new launch configuration, attach it to the existing Auto Scaling group, and terminate the running instances. Option A is invalid because Elastic beanstalk cannot launch new instances on demand. Since the current scenario requires Autoscaling, this is not the ideal option. Option B is invalid because this will be a maintenance overhead, since you just have an Autoscaling Group. There is no need to create a whole Cloudformation template for this. Option D is invalid because Autoscaling Group will still launch EC2 instances with the older launch configuration. For more information on Autoscaling Launch configuration, please refer to the below document link: from AWS http://docs.aws.amazon.com/autoscaling/latest/userguide/l_launchConfiguration.html

NEW QUESTION 3

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 4

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 5

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

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

Answer: C

Explanation:

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

NEW QUESTION 6

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

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

Answer: B

Explanation:

As your infrastructure grows, common patterns can emerge in which you declare the same components in each of your templates. You can separate out these common components and create dedicated templates for them. That way, you can mix and match different templates but use nested stacks to create a single, unified stack. Nested stacks are stacks that create other stacks. To create nested stacks, use the `AWS::CloudFormation::StackResource` in your template to reference other templates.

For more information on best practices for CloudFormation please refer to the below link: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html>

NEW QUESTION 7

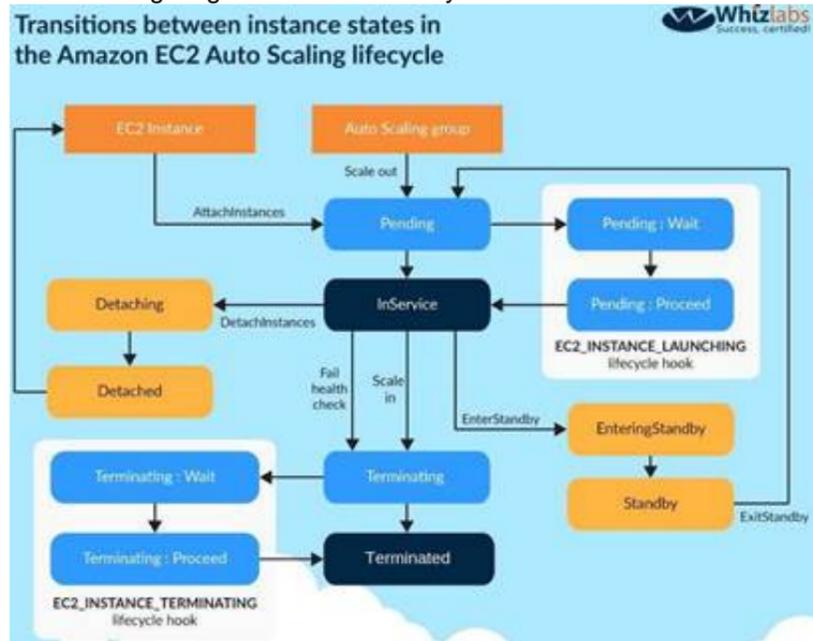
You have an Auto Scaling group of Instances that processes messages from an Amazon Simple Queue Service (SQS) queue. The group scales on the size of the queue. Processing Involves calling a third- party web service. The web service is complaining about the number of failed and repeated calls it is receiving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

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

Answer: D

Explanation:

The following diagram shows the lifecycle of the instances in Autoscaling



You can put the instances in a standby state, via the application, do the processing and then put the instance back in a state where it can be governed by the Autoscaling Group.

For more information on the Autoscaling Group Lifecycle please refer to the below link:

<http://docs.aws.amazon.com/autoscaling/latest/userguide/AutoScalingGroupLifecycle.html> | Note: As per AWS documentation.

To control whether an Auto Scaling group can terminate a particular instance when scaling in, use instance protection.

It is termed as Instance protection rather than termination protection when we refer it with "Scaling in process" of ASG.

For more information please view the following link: <https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-instance-termination.html#instance-protection-instance>

NEW QUESTION 8

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

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

Answer: CE

Explanation:

The AWS Security Token Service (STS) is a web service that enables you to request temporary, limited-privilege credentials for AWS Identity and Access Management (IAM) users or for users that you authenticate (federated users). The token can then be used to grant access to the objects in S3.

You can then provides access to the objects based on the key values generated via the user id. Option A is possible but then becomes a maintenance overhead because of the number of buckets. Option B is invalid because IAM users is not a good security practice.

Option D is invalid because SMS tokens are not efficient for this requirement.

For more information on the Security Token Service please refer to the below link: <http://docs.aws.amazon.com/STS/latest/APIReference/Welcome.html>

NEW QUESTION 9

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

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

Answer: A

Explanation:

As your infrastructure grows, common patterns can emerge in which you declare the same components in each of your templates. You can separate out these common components and create dedicated templates for them. That way, you can mix and match different templates but use nested stacks to create a single, unified stack. Nested stacks are stacks that create other stacks. To create nested stacks, use the `AWS::CloudFormation::StackResource` in your template to reference

other templates. For more information on CloudFormation best practises please refer to the below link:
<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html>

NEW QUESTION 10

You have an Auto Scaling group with 2 AZs. One AZ has 4 EC2 instances and the other has 3 EC2 instances. None of the instances are protected from scale in. Based on the default Auto Scaling termination policy what will happen?

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

Answer: D

Explanation:

The default termination policy is designed to help ensure that your network architecture spans Availability Zones evenly. When using the default termination policy, Auto Scaling selects an instance to terminate as follows:

Auto Scaling determines whether there are instances in multiple Availability Zones. If so, it selects the Availability Zone with the most instances and at least one instance that is not protected from scale in. If there is more than one Availability Zone with this number of instances, Auto Scaling selects the Availability Zone with the instances that use the oldest launch configuration. For more information on Autoscaling instance termination please refer to the below link:
<http://docs.aws.amazon.com/autoscaling/latest/userguide/as-instance-termination.html>

NEW QUESTION 10

If your application performs operations or workflows that take a long time to complete, what service can the Elastic Beanstalk environment do for you?

- A. Manages a Amazon SQS queue and running a daemon process on each instance
- B. Manages a Amazon SNS Topic and running a daemon process on each instance
- C. Manages Lambda functions and running a daemon process on each instance
- D. Manages the ELB and running a daemon process on each instance

Answer: A

Explanation:

Elastic Beanstalk simplifies this process by managing the Amazon SQS queue and running a daemon process on each instance that reads from the queue for you. When the daemon pulls an item from the queue, it sends an HTTP POST request locally to `http://localhost/` with the contents of the queue message in the body. All that your application needs to do is perform the long-running task in response to the POST.

For more information Elastic Beanstalk managing worker environments, please visit the below URL:
[? http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.html](http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.html)

NEW QUESTION 15

You are using Elastic Beanstalk to manage your e-commerce store. The store is based on an open source e-commerce platform and is deployed across multiple instances in an Auto Scaling group. Your development team often creates new "extensions" for the e-commerce store. These extensions include PHP source code as well as an SQL upgrade script used to make any necessary updates to the database schema. You have noticed that some extension deployments fail due to an error when running the SQL upgrade script. After further investigation, you realize that this is because the SQL script is being executed on all of your Amazon EC2 instances. How would you ensure that the SQL script is only executed once per deployment regardless of how many Amazon EC2 instances are running at the time?

- A. Use a "Container command" within an Elastic Beanstalk configuration file to execute the script, ensuring that the "leader only" flag is set to true.
- B. Make use of the Amazon EC2 metadata service to query whether the instance is marked as the leader" in the Auto Scaling group
- C. Only execute the script if "true" is returned.
- D. Use a "Solo Command" within an Elastic Beanstalk configuration file to execute the script
- E. The Elastic Beanstalk service will ensure that the command is only executed once.
- F. Update the Amazon RDS security group to only allow write access from a single instance in the Auto Scaling group; that way, only one instance will successfully execute the script on the database.

Answer: A

Explanation:

You can use the `container_commands` key to execute commands that affect your application source code. Container commands run after the application and web server have been set up and the application version archive has been extracted, but before the application version is deployed. Non-container commands and other customization operations are performed prior to the application source code being extracted.

You can use `leader_only` to only run the command on a single instance, or configure a test to only run the command when a test command evaluates to true. Leader-only container commands are only executed during environment creation and deployments, while other commands and server customization operations are performed every time an instance is provisioned or updated. Leader-only container commands are not executed due to launch configuration changes, such as a change in the AMI Id or instance type. For more information on customizing containers, please visit the below URL:
<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/customize-containers-ec2.html>

NEW QUESTION 16

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::StackResource` in your template to reference other templates.

For more information on CloudFormation best practices please refer to the below link: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html>

NEW QUESTION 20

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

- A. Deploy your web application as an AWS Elastic Beanstalk 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 23

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

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

Answer: AD

Explanation:

There is a rich brand of CLI commands available for EC2 Instances. The CLI is located in the following link:

- <http://docs.aws.amazon.com/cli/latest/reference/ec2/>

You can then use the `describe instances` command to describe the EC2 instances.

If you specify one or more instance IDs, Amazon EC2 returns information for those instances. If you do not specify instance IDs, Amazon EC2 returns information for all relevant instances. If you specify an instance ID that is not valid, an error is returned. If you specify an instance that you do not own, it is not included in the returned results.

- <http://docs.aws.amazon.com/cli/latest/reference/ec2/describe-instances.html>

You can use the `CC2 instances` to get those instances which need to be removed from the configuration management system.

NEW QUESTION 27

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 Dockers and Elastic beanstalk please refer to the below link:

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

NEW QUESTION 30

Your application uses CloudFormation to orchestrate your application's resources. During your testing phase before the application went live, your Amazon RDS instance type was changed and caused the instance to be re-created, resulting in the loss of test data. How should you prevent this from occurring in the future?

- A. Within the AWS CloudFormation parameter with which users can select the Amazon RDS instance type, set AllowedValues to only contain the current instance type.
- B. Use an AWS CloudFormation stack policy to deny updates to the instance.
- C. Only allow UpdateStack permission to IAM principals that are denied SetStackPolicy.
- D. In the AWS CloudFormation template, set the AWS::RDS::DBInstance's DBInstanceClass property to be read-only.
- E. Subscribe to the AWS CloudFormation notification "BeforeResourceUpdate," and call CancelStackUpdate if the resource identified is the Amazon RDS instance.
- F. Update the stack using ChangeSets

Answer: E

Explanation:

When you need to update a stack, understanding how your changes will affect running resources before you implement them can help you update stacks with confidence. Change sets allow you to preview how proposed changes to a stack might impact your running resources, for example, whether your changes will delete or replace any critical resources. AWS CloudFormation makes the changes to your stack only when you decide to execute the change set, allowing you to decide whether to proceed with your proposed changes or explore other changes by creating another change set.

For example, you can use a change set to verify that AWS CloudFormation won't replace your stack's database instances during an update.

NEW QUESTION 35

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.html>

NEW QUESTION 36

You currently have an Auto Scaling group with an Elastic Load Balancer and need to phase out all instances and replace with a new instance type. What are 2 ways in which this can be achieved.

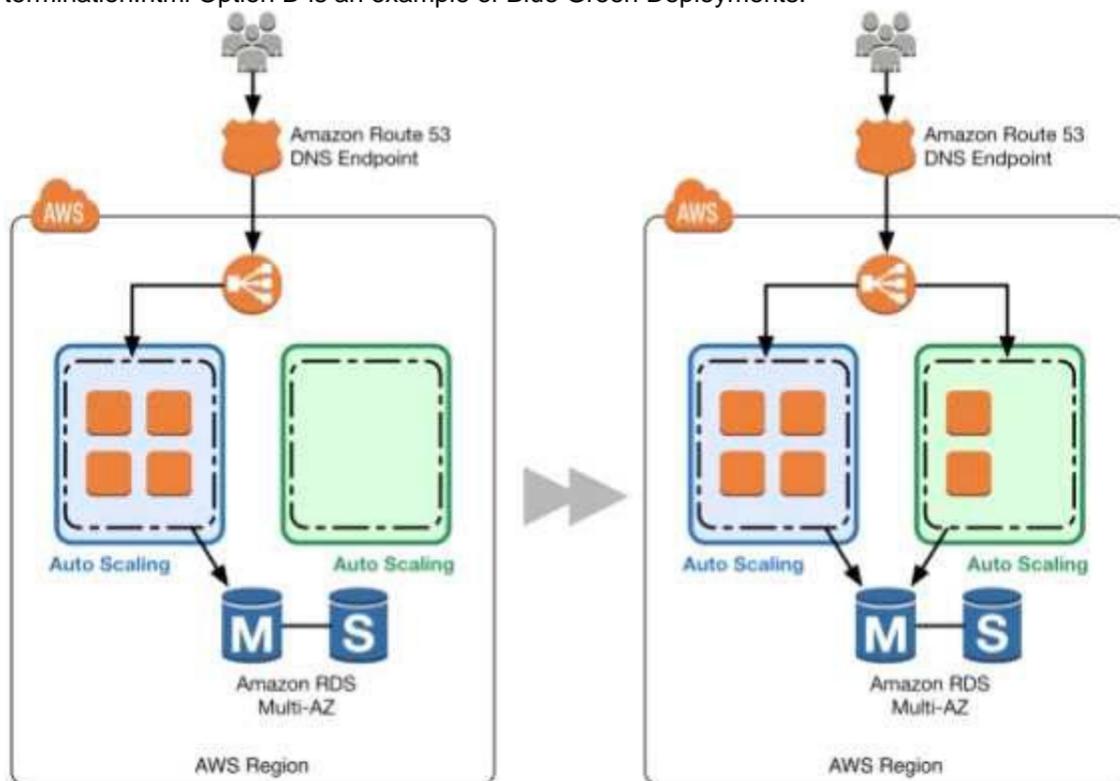
- A. Use Newest Instance to phase out all instances that use the previous configuration.
- B. Attach an additional ELB to your Auto Scaling configuration and phase in newer instances while removing older instances.
- C. Use OldestLaunchConfiguration to phase out all instances that use the previous configuration.
- D. V
- E. Attach an additional Auto Scaling configuration behind the ELB and phase in newer instances while removing older instances.

Answer: CD

Explanation:

When using the OldestLaunchConfiguration policy Auto Scaling terminates instances that have the oldest launch configuration. This policy is useful when you're updating a group and phasing out the instances from a previous configuration.

For more information on Autoscaling instance termination, please visit the below URL: <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-instance-termination.html> Option D is an example of Blue Green Deployments.



A blue group carries the production load while a green group is staged and deployed with the new code. When it's time to deploy, you simply attach the green group to the existing load balancer to introduce traffic to the new environment. For HTTP/HTTPS listeners, the load balancer favors the green Auto Scaling group because it uses a least outstanding requests routing algorithm.

As you scale up the green Auto Scaling group, you can take blue Auto Scaling group instances out of service by either terminating them or putting them in Standby state.

For more information on Blue Green Deployments, please refer to the below document link: from AWS

- https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 38

Your company has developed a web application and is hosting it in an Amazon S3 bucket configured for static website hosting. The application is using the AWS SDK for JavaScript in the browser to access data stored in an Amazon DynamoDB table. How can you ensure that API keys for access to your data in DynamoDB are kept secure?

- A. Create an Amazon S3 role in IAM with access to the specific DynamoDB tables, and assign it to the bucket hosting your website.
- B. Configure S3 bucket tags with your AWS access keys for your bucket hosting your website so that the application can query them for access.
- C. Configure a web identity federation role within IAM to enable access to the correct DynamoDB resources and retrieve temporary credentials.
- D. Store AWS keys in global variables within your application and configure the application to use these credentials when making requests.

Answer: C

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 document link: from AWS http://docs.wsamazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html

NEW QUESTION 42

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

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

Answer: B

Explanation:

Ensure first that the cloudformation template is updated with the new instance type. The AWS::AutoScaling::AutoScalingGroup resource supports an UpdatePolicy attribute. This is used to define how an Auto Scaling group resource is updated when an update to the Cloud Formation stack occurs. A common approach to updating an Auto Scaling group is to perform a rolling update, which is done by specifying the AutoScalingRollingUpdate policy. This retains the same Auto Scaling group and replaces old instances with new ones, according to the parameters specified. Option A is invalid because this will cause an interruption to the users.

Option C is partially correct, but it does not have all the steps as mentioned in option B.
Option D is partially correct, but we need the AutoScalingRollingUpdate attribute to ensure a rolling update is performed.
For more information on AutoScaling Rolling updates please refer to the below link:
• <https://aws.amazon.com/premiumsupport/knowledge-center/auto-scaling-group-rolling-updates/>

NEW QUESTION 47

You need to create a Route53 record automatically in CloudFormation when not running in production during all launches of a Template. How should you implement this?

- A. Use a Parameter for environment, and add a Condition on the Route53 Resource in the template to create the record only when environment is not production.
- B. Create two templates, one with the Route53 record value and one with a null value for the record.
- C. Use the one without it when deploying to production.
- D. Use a Parameter for environment, and add a Condition on the Route53 Resource in the template to create the record with a null string when environment is production.
- E. Create two templates, one with the Route53 record and one without it.
- F. Use the one without it when deploying to production.

Answer: A

Explanation:

The optional Conditions section includes statements that define when a resource is created or when a property is defined. For example, you can compare whether a value is equal to another value. Based on the result of that condition, you can conditionally create resources. If you have multiple conditions, separate them with commas.

You might use conditions when you want to reuse a template that can create resources in different contexts, such as a test environment versus a production environment. In your template, you can add an EnvironmentType input parameter, which accepts either prod or test as inputs. For the production environment, you might include Amazon EC2 instances with certain capabilities; however, for the test environment, you want to use reduced capabilities to save money. With conditions, you can define which resources are created and how they're configured for each environment type.

For more information on CloudFormation conditions please refer to the below link: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/conditions-section-structure.html>

NEW QUESTION 49

You have a development team that is continuously spending a lot of time rolling back updates for an application. They work on changes, and if the change fails, they spend more than 5-6h in rolling back the update. Which of the below options can help reduce the time for rolling back application versions.

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

Answer: A

Explanation:

Option B is invalid because Elastic Beanstalk already has the facility to manage various versions and you don't need to use S3 separately for this.

Option C is invalid because in CloudFormation you will need to maintain the versions. Elastic Beanstalk can do that automatically for you.

Option D is good for production scenarios and Elastic Beanstalk is great for development scenarios. AWS Beanstalk is the perfect solution for developers to maintain application versions.

With AWS Elastic Beanstalk, you can quickly deploy and manage applications in the AWS Cloud without worrying about the infrastructure that runs those applications. AWS Elastic Beanstalk reduces management complexity without restricting choice or control. You simply upload your application, and AWS Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring.

For more information on AWS Beanstalk please refer to the below link: <https://aws.amazon.com/documentation/elastic-beanstalk/>

NEW QUESTION 50

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

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

Answer: C

Explanation:

The first step in using Elastic Beanstalk is to create an application, which represents your web application in AWS. In Elastic Beanstalk an application serves as a container for the environments that run your web app, and versions of your web app's source code, saved configurations, logs and other artifacts that you create while using Elastic Beanstalk.

For more information on Applications, please refer to the below link: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/applications.html>

Deploying a new version of your application to an environment is typically a fairly quick process. The new source bundle is deployed to an instance and extracted, and the web container or application server picks up the new version and restarts if necessary. During deployment, your application might still become unavailable to users for a few seconds. You can prevent this by configuring your environment to use rolling deployments to deploy the new version to instances in batches. For more information on deployment, please refer to the below link: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-deploy-existing-version.html>

NEW QUESTION 53

You have decided to migrate your application to the cloud. You cannot afford any downtime. You want to gradually migrate so that you can test the application with a small percentage of users and increase over time. Which of these options should you implement?

- A. Use Direct Connect to route traffic to the on-premise location.
- B. In DirectConnect, configure the amount of traffic to be routed to the on-premise location.
- C. Implement a Route 53 failover routing policy that sends traffic back to the on-premise application if the AWS application fails.

- D. Configure an Elastic Load Balancer to distribute the traffic between the on-premises application and the AWS application.
- E. Implement a Route 53 weighted routing policy that distributes the traffic between your on-premises application and the AWS application depending on weight.

Answer: D

Explanation:

Option A is incorrect because DirectConnect cannot control the flow of traffic.

Option B is incorrect because you want to split the percentage of traffic. Failover will direct all of the traffic to the backup servers.

Option C is incorrect because you cannot control the percentage distribution of traffic.

Weighted routing lets you associate multiple resources with a single domain name (example.com) or subdomain name (acme.example.com) and choose how much traffic is routed to each resource. This can be useful for a variety of purposes, including load balancing and testing new versions of software.

For more information on the Routing policy please refer to the below link: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

NEW QUESTION 55

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 clusterjobs to perform adhoc MapReduce analysis and write new queries when needed.
- G. Copy all log files into AWS S3 using a cron job on each 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 Elastic search 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.

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 Cloud Watch so that you can go from raw data to actionable insights quickly. For more information on Elastic Search, please refer to the below link:

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

NEW QUESTION 58

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

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

Answer: ACD

Explanation:

You can use Cloud Watch Logs to monitor applications and systems using log data. For example,

CloudWatch Logs can track the number of errors that occur in your

application logs and send you a notification whenever the rate of errors exceeds a threshold you specify. CloudWatch Logs uses your log data for monitoring; so, no code changes are required. For example, you can monitor application logs for specific literal terms (such as "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, CloudWatch Logs reports the data to a CloudWatch metric that you specify. For more information on Cloudwatch Logs please refer to the below link:

<http://docs.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 Cloudwatch and SNS please refer to the below link:

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

NEW QUESTION 61

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

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

Answer: C

Explanation:

Answer - C

You can use Kinesis Streams for rapid and continuous data intake and aggregation. The type of data used includes IT infrastructure log data, application logs, social media, market data feeds, and web clickstream data. Because the response time for the data intake and processing is in real time, the processing is typically lightweight.

The following are typical scenarios for using Kinesis Streams:

- Accelerated log and data feed intake and processing - You can have producers push data directly into a stream. For example, push system and application logs and they'll be available for processing in seconds. This prevents the log data from being lost if the front end or application server fails. Kinesis Streams provides accelerated data feed intake because you don't batch the data on the servers before you submit it for intake.
 - Real-time metrics and reporting - You can use data collected into Kinesis Streams for simple data analysis and reporting in real time. For example, your data-processing application can work on metrics and reporting for system and application logs as the data is streaming in, rather than wait to receive batches of data.
- For more information on Amazon Kinesis and SNS please refer to the below link:
- <http://docs.aws.amazon.com/streams/latest/dev/introduction.html>

NEW QUESTION 66

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

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

Answer: C

Explanation:

You can use AWS CloudTrail to get a history of AWS API calls and related events for your account. This history includes calls made with the AWS Management Console, AWS Command Line Interface, AWS SDKs, and other AWS services. For more information on Cloudtrail, please visit the below URL:

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

Amazon Cloud Watch Events 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 Events becomes aware of operational changes as they occur. Cloud Watch Events responds to these operational changes and takes corrective action as necessary, by sending messages to respond to the environment, activating functions, making changes, and capturing state information. For more information on Cloud watch events, please visit the below URL:

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

NEW QUESTION 71

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

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

Answer: C

Explanation:

The most likely cause is the fact you have hit the maximum of 5 Elastic IP's per region.

By default, all AWS accounts are limited to 5 Elastic IP addresses per region, because public (IPv4) Internet addresses are a scarce public resource. We strongly encourage you to use an Elastic IP address primarily for the ability to remap the address to another instance in the case of instance failure, and to use DNS hostnames for all other inter-node communication.

Option A is invalid because a AMI does not have a Development version tag. Option B is invalid because there is no flag for an EC2 Instance

Option D is invalid because there is a limit of 5 VPCs per region. For more information on Elastic IP's, please visit the below URL:

- <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html>

NEW QUESTION 76

You want to pass queue messages that are 1GB each. How should you achieve 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. Use SQS's support for message partitioning and multi-part uploads on Amazon S3.
- E. Use AWS EFS as a shared pool storage medium
- F. Store filesystem pointers to the files on disk in the SQS message bodies.

Answer: B

Explanation:

You can manage Amazon SQS messages with Amazon S3. This is especially useful for storing and consuming messages with a message size of up to 2 GB. To manage

Amazon SQS messages with Amazon S3, use the Amazon SQS Extended Client Library for Java. Specifically, you use this library to:

- Specify whether messages are always stored in Amazon S3 or only when a message's size exceeds 256 KB.
- Send a message that references a single message object stored in an Amazon S3 bucket.
- Get the corresponding message object from an Amazon S3 bucket.
- Delete the corresponding message object from an Amazon S3 bucket.

For more information on processing large messages for SQS, please visit the below URL:

- <http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-s3-messages.html>

NEW QUESTION 81

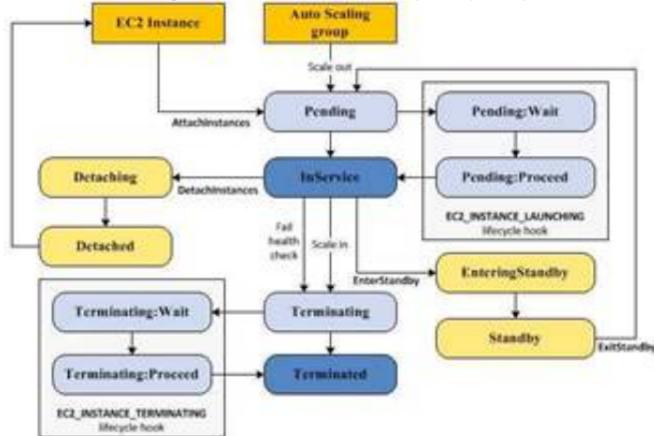
For AWS Auto Scaling, what is the first transition state an existing instance enters after leaving Standby state?

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

Answer: C

Explanation:

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



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

NEW QUESTION 82

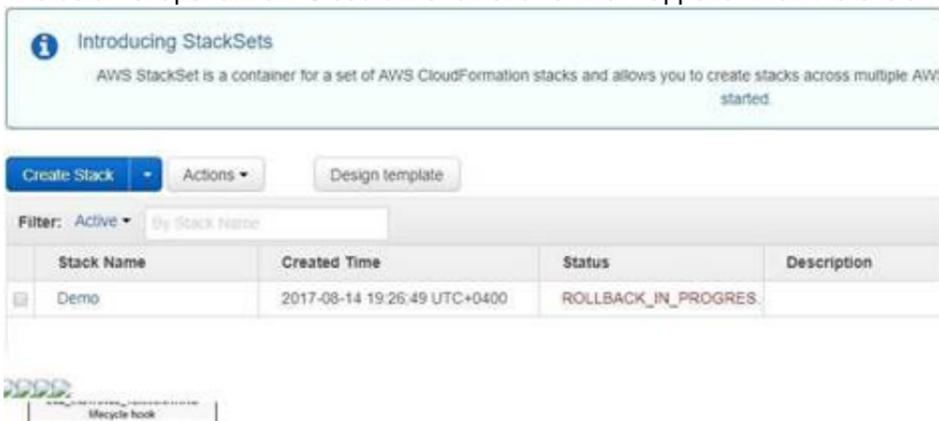
You have deployed a Cloudformation template which is used to spin up resources in your account. Which of the following status in Cloudformation represents a failure.

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

Answer: C

Explanation:

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



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

NEW QUESTION 85

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

- A. io1
- B. gp1
- C. standard
- D. gp2

Answer: A

Explanation:

Provisioned IOPS SSD should be used for critical business applications that require sustained IOPS performance, or more than 10,000 IOPS or 160 MiB/s of throughput per volume
 This is ideal for Large database workloads, such as:

- MongoDB
- Cassandra
- MicrosoftSQL Server
- MySQL
- PostgreSQL

- Oracle
- For more information on the various CBS Volume Types, please refer to the below link:
- <http://docs.aws.amazon.com/AWSSCC2/latest/UserGuide/CBSVolumeTypes.html>

NEW QUESTION 89

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

- A. Purchase a Heavy Utilization Reserved Instance to run the accounting softwar
- B. Turn it off after hour
- C. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.
- D. Purchase a Medium Utilization Reserved Instance to run the accounting softwar
- E. Turn it off after hour
- F. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.
- G. Purchase a Light Utilization Reserved Instance to run the accounting softwar
- H. Turn it off after hour
- I. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.
- J. Purchase a Full Utilization Reserved Instance to run the accounting softwar
- K. Turn it off after hour
- L. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.

Answer: A

Explanation:

Reserved Instances provide you with a significant discount compared to On-Demand Instance pricing.

Reserved Instances are not physical instances, but rather a

billing discount applied to the use of On-Demand Instances in your account. These On-Demand Instances must match certain attributes in order to benefit from the billing discount

For more information, please refer to the below link:

- <https://aws.amazon.com/about-aws/whats-new/2011/12/01/New-Amazon-CC2-Reserved-Instances-Options-Now-Available/>
- <https://aws.amazon.com/blogs/aws/reserved-instance-options-for-amazon-ec2/>
- <http://docs.aws.amazon.com/AWSSCC2/latest/UserGuide/ec2-reserved-instances.html> Note:

It looks like these options are also no more available at present.

It looks like Convertible, Standard and scheduled are the new instance options. However the exams may still be referring to the old RIs.

<https://aws.amazon.com/ec2/pricing/reserved-instances/>

NEW QUESTION 90

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

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

Answer: A

Explanation:

Amazon Route 53 health checks monitor the health and performance of your web applications, web servers, and other resources. Each health check that you create

can monitor one of the following:

- The health of a specified resource, such as a web server
- The status of an Amazon Cloud Watch alarm
- The status of other health checks

For more information on Route53 Health checks, please refer to the below link:

- <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html>

NEW QUESTION 94

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

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

Answer: C

Explanation:

Amazon RDS Read Replicas provide enhanced performance and durability for database (DB) instances. This replication feature makes it easy to elastically scale out beyond the capacity constraints of a single DB Instance for read-heavy database workloads. You can create one or more replicas of a given source DB Instance and serve high-volume application read traffic from multiple copies of your data, thereby increasing aggregate read throughput. Read replicas can also be promoted when needed to become standalone DB instances.

Option A is invalid because you would need to maintain the synchronization yourself with a secondary instance.

Option B is invalid because you are introducing another layer unnecessarily when you already have read replica's Option D is invalid because you only use this for Standby's

For more information on Read Replica's, please refer to the below link: <https://aws.amazon.com/rds/details/read-replicas/>

NEW QUESTION 99

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, 1AM Credential Reports
- B. CloudTrail, 1AM Credential Reports, AWS Config
- C. CloudTrail, AWS Config, 1AM Credential Reports
- D. AWS Config, 1AM 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/awsccloudtrail/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 1AM 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 104

You are building a game high score table in DynamoDB. You will store each user's highest score for each game, with many games, all of which have relatively similar usage levels and numbers of players. You need to be able to look up the highest score for any game. What's the best DynamoDB key structure?

- A. HighestScore as the hash/only key.
- B. GameID as the hash key, HighestScore as the range key
- C. GameID as the hash/only key.
- D. GameID as the hash/only key.

Answer: B

Explanation:

It always best to choose the hash key as the column that will have a wide range of values. This is also given in the AWS documentation

Choosing a Partition Key

The following table compares some common partition key schemas for provisioned throughput efficiency:

Partition key value	Uniformity
User ID, where the application has many users.	Good
Status code, where there are only a few possible status codes.	Bad
Item creation date, rounded to the nearest time period (e.g. day, hour, minute)	Bad
Device ID, where each device accesses data at relatively similar intervals	Good
Device ID, where even if there are a lot of devices being tracked, one is by far more popular than all the others.	Bad

Next since you need to sort by the Highest Score, you need to use that as the sort key For more information on Table Guidelines, please visit the below URL:

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

NEW QUESTION 107

You need to perform ad-hoc business analytics queries on well-structured data. Data comes in constantly at a high velocity. Your business intelligence team can understand SQL.

What AWS service(s) should you look to first?

- A. Kinesis Firehose + RDS
- B. Kinesis Firehose+RedShift
- C. EMR using Hive
- D. EMR running Apache Spark

Answer: B

Explanation:

Amazon Kinesis Firehose is the easiest way to load streaming data into AWS. It can capture, transform, and load streaming data into Amazon Kinesis Analytics, Amazon S3, Amazon Redshift, and Amazon Elasticsearch Service, enabling near real-time analytics with existing business intelligence tools and dashboards you're already using today. It is a fully managed service that automatically scales to match the throughput of your data and requires no ongoing administration. It can also batch, compress, and encrypt the data before loading it, minimizing the amount of storage used at the destination and increasing security.

For more information on Kinesis firehose, please visit the below URL:

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

Amazon Redshift is a fully managed, petabyte-scale data warehouse service in the cloud. You can start with just a few hundred gigabytes of data and scale to a petabyte or more. This enables you to use your data to acquire new insights for your business and customers. For more information on Redshift, please visit the below URL:

<http://docs.aws.amazon.com/redshift/latest/mgmt/welcome.html>

NEW QUESTION 111

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 115

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

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

Answer: B

Explanation:

With AWS Cloud Formation, you can update the properties for resources in your existing stacks.

These changes can range from simple configuration changes, such

as updating the alarm threshold on a Cloud Watch alarm, to more complex changes, such as updating the Amazon Machine Image (AMI) running on an Amazon EC2

instance. Many of the AWS resources in a template can be updated, and we continue to add support for more.

For more information on Cloudformation version control, please visit the below URL:

http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/updates_stacks.html

NEW QUESTION 119

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

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

Answer: A

Explanation:

This question is more on the grounds of validating each option

Option B is invalid, because the Route table would have an effect on all worker processes and no jobs would have been completed.

Option C is invalid because if the IAM Role was invalid then no jobs would be completed.

Option D is invalid because the scaling is happening, its just that the jobs are not getting completed. For more information on Scaling on Demand, please visit the below URL:

• <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scale-based-on-demand.html>

NEW QUESTION 124

If I want Cloud Formation stack status updates to show up in a continuous delivery system in as close to real time as possible, how should I achieve this?

- A. Use a long-poll on the Resources object in your Cloud Formation stack and display those state changes in the UI for the system.
- B. Use a long-poll on the ListStacksAPI call for your CloudFormation stack and display those state changes in the UI for the system.
- C. Subscribe your continuous delivery system to an SNS topic that you also tell your CloudFormation stack to publish events into.
- D. Subscribe your continuous delivery system to an SQS queue that you also tell your CloudFormation stack to publish events into.

Answer: C

Explanation:

Answer - C

You can monitor the progress of a stack update by viewing the stack's events. The console's Cvents tab displays each major step in the creation and update of the stack sorted by the time of each event with latest events on top. The start of the stack update process is marked with an UPDATE_IN_PROGRESS event for the

stack For more information on Monitoring your stack, please visit the below URL:
<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks-monitor-stack.html>

NEW QUESTION 125

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

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

Answer: D

Explanation:

A placement group is a logical grouping of instances within a single Availability Zone. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. To provide the lowest latency, and the highest packet-per-second network performance for your placement group, choose an instance type that supports enhanced networking. For more information on Placement Groups, please visit the below URL:
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

NEW QUESTION 126

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 131

You are building a mobile app for consumers to post cat pictures online. You will be storing the images in AWS S3. You want to run the system very cheaply and simply. Which one of these options allows you to build a photo sharing application with the right authentication/authorization implementation.

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

Answer: A

Explanation:

Amazon Cognito lets you easily add user sign-up and sign-in and manage permissions for your mobile and web apps. You can create your own user directory within Amazon Cognito. You can also choose to authenticate users through social identity providers such as Facebook, Twitter, or Amazon; with SAML identity solutions; or by using your own identity system. In addition, Amazon Cognito enables you to save data locally on users' devices, allowing your applications to work even when the devices are offline. You can then synchronize data across users' devices so that their app experience remains consistent regardless of the device they use.

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

- <http://docs.aws.amazon.com/cognito/latest/developerguide/what-is-amazon-cognito.html>

NEW QUESTION 136

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

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

Answer: B

Explanation:

AD Connector is a directory gateway with which you can redirect directory requests to your on-premises Microsoft Active Directory without caching any information in the cloud. AD Connector comes in two sizes, small and large. A small AD Connector is designed for smaller organizations of up to 500 users. A large AD Connector

can support larger organizations of up to 5,000 users. Once set up, AD Connector offers the following benefits:

- Your end users and IT administrators can use their existing corporate credentials to log on to AWS applications such as Amazon Workspaces, Amazon WorkDocs, or Amazon WorkMail.
- You can manage AWS resources like Amazon EC2 instances or Amazon S3 buckets through IAM role-based access to the AWS Management Console.
- You can consistently enforce existing security policies (such as password expiration, password history, and account lockouts) whether users or IT administrators are accessing resources in your on-premises infrastructure or in the AWS Cloud.
- You can use AD Connector to enable multi-factor authentication by integrating with your existing RADIUS-based MFA infrastructure to provide an additional layer of security when users access AWS applications.

For more information on the AD Connector, please visit the below URL:

- http://docs.aws.amazon.com/directoryservice/latest/admin-guide/directory_ad_connector.htm

NEW QUESTION 141

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

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

Answer: D

Explanation:

Custom resources enable you to write custom provisioning logic in templates that AWS CloudFormation runs anytime you create, update (if you changed the custom resource), or delete stacks. For example, you might want to include resources that aren't available as AWS CloudFormation resource types. You can include those resources by using custom resources. That way you can still manage all your related resources in a single stack.

Use the AWS::CloudFormation::CustomResource or Custom::String resource type to define custom resources in your templates. Custom resources require one property: the service token, which specifies where AWS CloudFormation sends requests to, such as an Amazon SNS topic.

For more information on Custom Resources in CloudFormation, please visit the below URL:

- <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-custom-resources.html>

NEW QUESTION 145

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 150

You need to grant a vendor access to your AWS account. They need to be able to read protected messages in a private S3 bucket at their leisure. They also use AWS. What is the best way to accomplish this?

- A. Create an IAM User with API Access Key
- B. Grant the User permissions to access the bucket
- C. Give the vendor the AWS Access Key ID and AWS Secret Access Key for the User.
- D. Create an EC2 Instance Profile on your account
- E. Grant the associated IAM role full access to the bucket
- F. Start an EC2 instance with this Profile and give SSH access to the instance to the vendor.
- G. Create a cross-account IAM Role with permission to access the bucket, and grant permission to use the Role to the vendor AWS account.
- D- Generate a signed S3 PUT URL and a signed S3 GET URL, both with wildcard values and 2 year duration
- H. Pass the URLs to the vendor.

Answer: C

Explanation:

You can use AWS Identity and Access Management (IAM) roles and AWS Security Token Service (STS) to set up cross-account access between AWS accounts.

When you assume an IAM role in another AWS account to obtain cross-account access to services and resources in that account, AWS CloudTrail logs the cross-account activity. For more information on Cross Account Access, please visit the below URL:

- <https://aws.amazon.com/blogs/security/tag/cross-account-access/>

NEW QUESTION 151

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 created
- C. If older than 2 months, delete them and recreate new keys
- D. Delete the user associated with the keys after every 2 months
- E. Then recreate the user again.
- D- Delete the IAM Role associated with the keys after every 2 months
- F. Then recreate the IAM Role again.

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 `list-access-keys` CLI command returns information about the access key IDs associated with the specified IAM 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 155

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. Review the 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 occurring 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 160

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

- 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 163

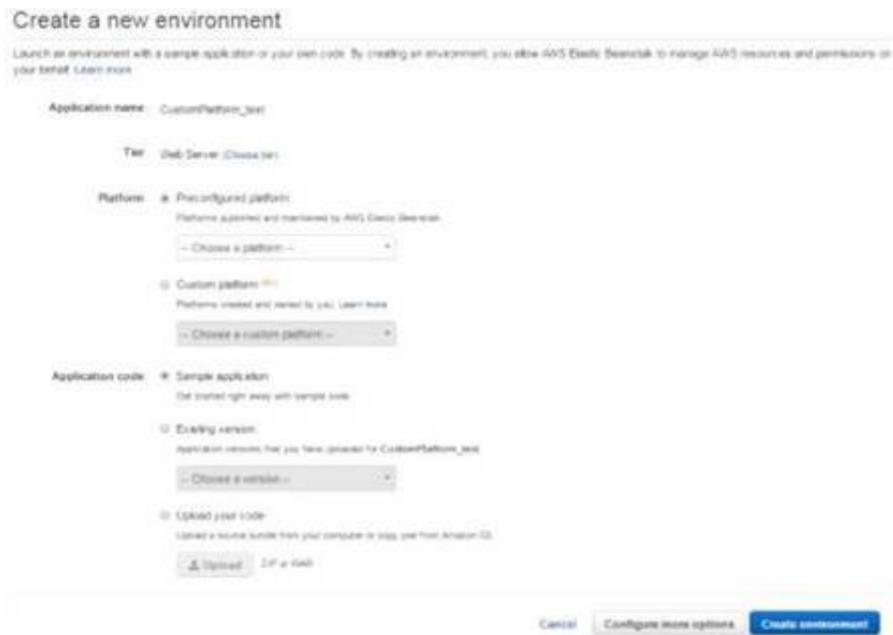
When creating an Elastic Beanstalk environment using the Wizard, what are the 3 configuration options presented to you?

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

Answer: ABD

Explanation:

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



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

NEW QUESTION 167

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

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

Answer: B

Explanation:

The AWS Documentation mentions

The load balancer routes requests only to the healthy instances. When the load balancer determines that an instance is unhealthy, it stops routing requests to that instance. The load balancer resumes routing requests to the instance when it has been restored to a healthy state.

For more information on ELB health checks, please refer to the below link: <http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-healthchecks.html>

NEW QUESTION 171

Which of the following services can be used in conjunction with Cloudwatch Logs. Choose the 3 most viable services from the options given below

- A. Amazon Kinesis
- B. Amazon S3
- C. Amazon SQS
- D. Amazon Lambda

Answer: ABD

Explanation:

The AWS Documentation the following products which can be integrated with Cloudwatch logs

- 1) Amazon Kinesis - Here data can be fed for real time analysis
 - 2) Amazon S3 - You can use CloudWatch Logs to store your log data in highly durable storage such as S3.
 - 3) Amazon Lambda - Lambda functions can be designed to work with Cloudwatch log
- For more information on Cloudwatch Logs, please refer to the below link: [link:http://docs^ws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html](http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html)

NEW QUESTION 175

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

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

Answer: C

Explanation:

The AWS Documentation mentions

Because Elastic Beanstalk uses a drop-in upgrade process, there might be a few seconds of downtime. Use rolling deployments to minimize the effect of deployments on your production environments.

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

- <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/troubleshooting-deployments.html>
- <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.deploy-existing-version.html>

NEW QUESTION 176

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 181

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

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

Answer: B

Explanation:

The AWS documentation mentions

AWS CodeDeploy is a deployment service that automates application deployments to Amazon EC2 instances or on-premises instances in your own facility.

For more information on Code Deploy please refer to the below link:

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

NEW QUESTION 184

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

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

Answer: AD

Explanation:

The following deployment types are available

1. In-place deployment: The application on each instance in the deployment group is stopped, the latest application revision is installed, and the new version of the application is started and validated.

2. Blue/green deployment: The instances in a deployment group (the original environment) are replaced by a different set of instances (the replacement environment)

For more information on Code Deploy please refer to the below link:

- <http://docs.aws.amazon.com/codedeploy/latest/userguide/primary-components.html>

NEW QUESTION 188

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

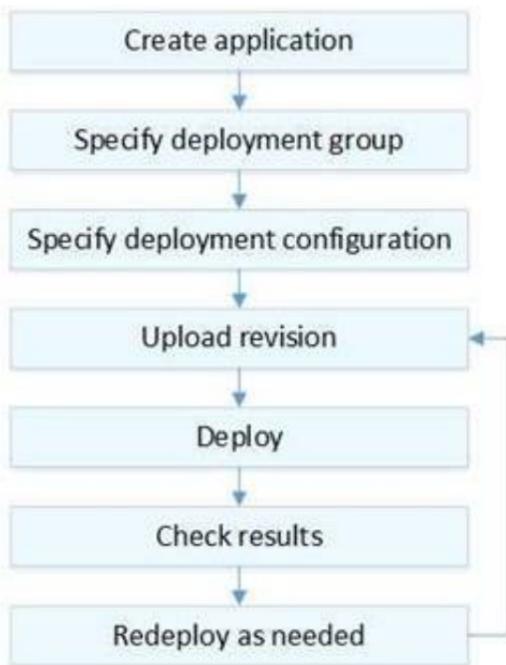
- 1) Specify deployment configuration
- 2) Upload revision
- 3) Create application
- 4) Specify deployment group

- 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 192

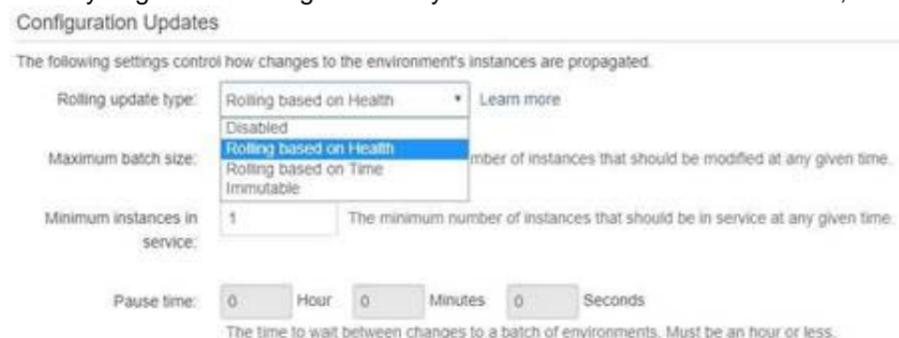
Which of the following is not a rolling type update which is present for Configuration Updates when it comes to the Elastic Beanstalk service

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

Answer: B

Explanation:

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



The AWS Documentation mentions

- 1) With health-based rolling updates. Elastic Beanstalk waits until instances in a batch pass health checks before moving on to the next batch.
- 2) For time-based rolling updates, you can configure the amount of time that Elastic Beanstalk waits after completing the launch of a batch of instances before moving on to the next batch. This pause time allows your application to 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 Scaling group. 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.rollingupdates.html>

NEW QUESTION 195

You have the requirement to get a snapshot of the current configuration of the resources in your AWS Account. Which of the following services can be used for this purpose

- A. AWS CodeDeploy
- B. AWS Trusted Advisor
- C. AWSConfig
- D. AWSIAM

Answer: C

Explanation:

The AWS Documentation mentions the following With AWS Config, you can do the following:

- Evaluate your AWS resource configurations for desired settings.
- Get a snapshot of the current configurations of the supported resources that are associated with your AWS account.
- Retrieve configurations of one or more resources that exist in your account.
- Retrieve historical configurations of one or more resources.
- Receive a notification whenever a resource is created, modified, or deleted.
- View relationships between resources. For example, you might want to find all resources that use a particular security group. For more information on AWS Config, please visit the below URL: <http://docs.aws.amazon.com/config/latest/developerguide/WhatIsConfig.html>

NEW QUESTION 199

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

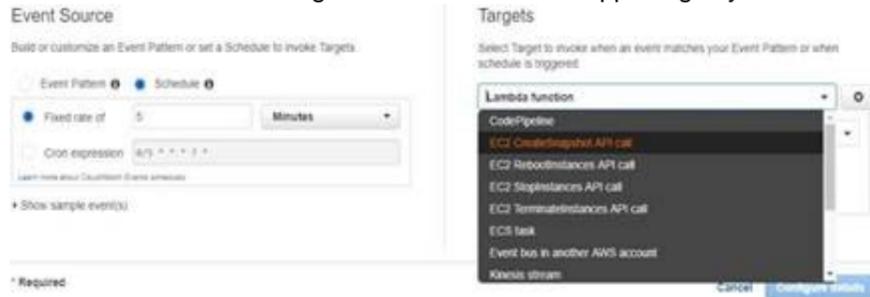
Answer: D

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 powershell script on Windows machines and maintaining the script itself. And then you have the overhead of having a separate instance just to run that script.

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

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



The AWS Documentation mentions

Amazon Cloud Watch Events 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 Events becomes aware of operational changes as they occur. Cloud Watch Events responds to these operational changes and takes corrective action as necessary, by sending messages to respond to the environment, activating functions, making changes, and capturing state information. For more information on Cloud watch Events, please visit the below URL:

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

NEW QUESTION 203

Your company has a set of resources hosted in AWS. They want to be notified when the costs of the AWS resources running in the account reaches a certain threshold. How can this be accomplished in an ideal way.

- A. Create a script which monitors all the running resources and calculates the costs accordingly.
- B. Download the cost reports and analyze the reports to see if the costs are going beyond the threshold
- C. Create a billing alarm which can alert you when the costs are going beyond a certain threshold
- D. Create a consolidated billing report and see if the costs are going beyond the threshold.

Answer: C

Explanation:

The AWS Documentation mentions

You can monitor your AWS costs by using Cloud Watch. With Cloud Watch, you can create billing alerts that notify you when your usage of your services exceeds thresholds that you define. You specify these threshold amounts when you create the billing alerts.

When your usage exceeds these amounts, AWS sends you an email notification. You can also sign up to receive notifications when AWS prices change. For more information on billing alarms, please visit the below URL:

- <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/monitor-charges.html>

NEW QUESTION 205

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 accordingly
- 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

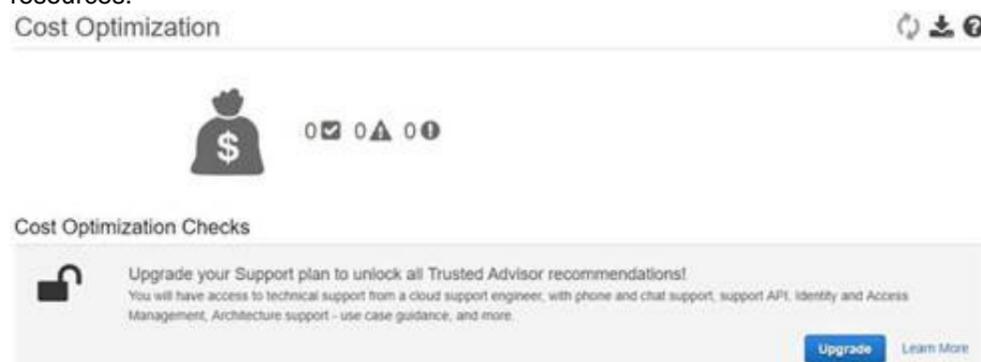
Explanation:

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 210

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 screenshot shows the AWS IAM console for a user named 'testuser'. Under 'Permissions summary', a table lists the following policies attached to the user:

Type	Name
Managed policy	AWSCloudTrailReadOnlyAccess
Managed policy	IAMUserChangePassword

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 212

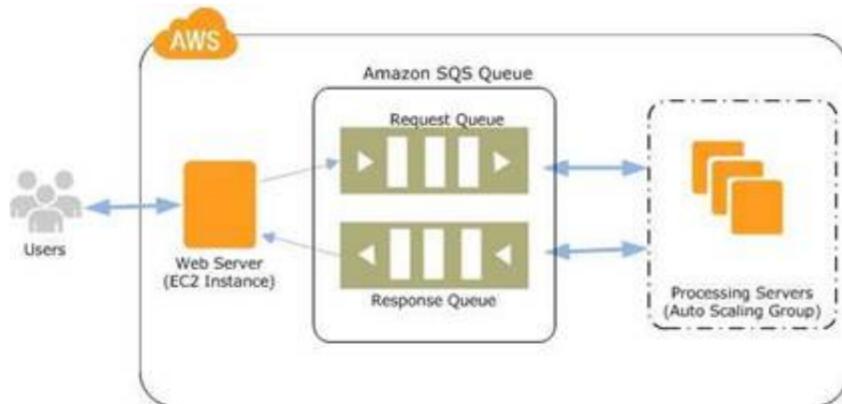
Which of the following design strategies is ideal when designing loosely coupled systems. Choose 2 answers from the options given below

- A. Having the web and worker roles running on the same set of EC2 Instances
- B. Having the web and worker roles running on separate EC2 Instances
- C. Using SNS to establish communication between the web and worker roles
- D. Using SQS to establish communication between the web and worker roles

Answer: BD

Explanation:

The below diagram shows the ideal design which uses SQS and separate environments for web and worker processes. The SQS queue manages the communication between the web and worker roles.



One example is the way Elastic beanstalk manages worker environments. For more information on this, please visit the below URL:

- ? <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.html>

NEW QUESTION 216

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 221

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

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

Answer: C

Explanation:

The AWS Documentation mentions the following

AD Connector is a directory gateway with which you can redirect directory requests to your on-premises Microsoft Active Directory without caching any information in the cloud. AD Connector comes in two sizes, small and large. A small AD Connector is designed for smaller organizations of up to 500 users. A large AD Connector can support larger organizations of up to 5,000 users.

For more information on the AD connector, please refer to the below URL: http://docs.aws.amazon.com/directoryservice/latest/admin-guide/directory_ad_connector.html

NEW QUESTION 226

You're building a mobile application game. The application needs permissions for each user to communicate and store data in DynamoDB tables. What is the best method for granting each mobile device that installs your application to access DynamoDB tables for storage when required? Choose the correct answer from the options below

- A. During the install and game configuration process, have each user create an IAM credential and assign the IAM user to a group with proper permissions to communicate with DynamoDB.
- B. Create an IAM group that only gives access to your application and to the DynamoDB table
- C. Then, when writing to DynamoDB, simply include the unique device ID to associate the data with that specific user.
- D. Create an IAM role with the proper permission policy to communicate with the DynamoDB table
- E. Use web identity federation, which assumes the IAM role using AssumeRoleWithWebIdentity, when the user signs in, granting temporary security credentials using STS.
- F. Create an Active Directory server and an AD user for each mobile application use
- G. When the user signs in to the AD sign-on, allow the AD server to federate using SAML 2.0 to IAM and assign a role to the AD user which is the assumed with AssumeRoleWithSAML

Answer: C

Explanation:

Answer - C

For access to any AWS service, the ideal approach for any application is to use Roles. This is the first preference.

For more information on IAM policies please refer to the below link:

http://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html

Next for any web application, you need to use web identity federation. Hence option D is the right option. This along with the usage of roles is highly stressed in the AWS documentation.

The AWS documentation mentions the following

When developing a web application it is recommended not to embed or distribute long-term AWS credentials with apps that a user downloads to a device, even in an encrypted store. Instead, build your app so that it requests temporary AWS security credentials dynamically when needed using web identity federation. The supplied temporary credentials map to an AWS role that has only the permissions needed to perform the tasks required by the mobile app.

For more information on web identity federation please refer to the below link: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html

NEW QUESTION 228

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 229

You are the IT administrator for your company. You have the responsibility of creating development environments which would conform to the LAMP development

stack. The requirement is that the development team always gets the latest version of the LAMP stack each time a new instance is launched. Which of the following is an efficient and effective way to implement this requirement? Choose 2 answers from the options given below

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

Answer: BC

Explanation:

Using User data and cloud-init directives you can always ensure you download the latest version of the LAMP stack and give it to the development teams. With AMI's

you will always have the same version and will need to create an AMI everytime the version of the LAMP stack changes.

The AWS Documentation mentions

When you launch an instance in Amazon 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 User data please refer to the below link: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/user-data.html>

NEW QUESTION 232

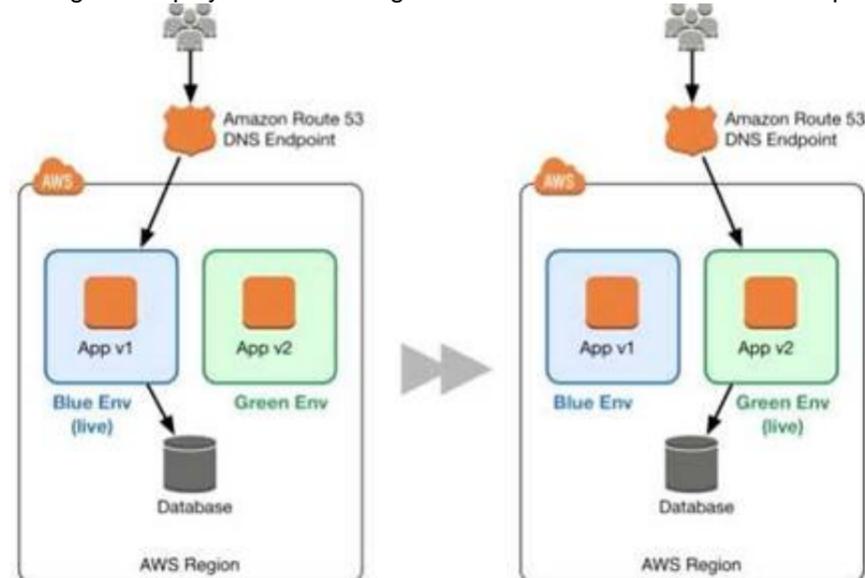
A company has developed a Ruby on Rails content management platform. Currently, OpsWorks with several stacks for dev, staging, and production is being used to deploy and manage the application. Now the company wants to start using Python instead of Ruby. How should the company manage the new deployment?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

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

Answer: B

Explanation:

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



Please find the below link on a white paper for blue green deployments

- https://d03wsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 233

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

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

Answer: C

Explanation:

The AWS documentation mentions the following

With the Deletion Policy attribute you can preserve or (in some cases) backup a resource when its stack is deleted. You specify a DeletionPolicy attribute for each resource that you want to control. If a resource has no DeletionPolicy attribute, AWS Cloud Formation deletes the resource by default. Note that this capability also applies to update operations that lead to resources being removed.

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

NEW QUESTION 238

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 242

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

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

Answer: B

Explanation:

The AWS documentation mentions

The Auto Scaling cooldown period is a configurable setting for your Auto Scaling group that helps to ensure that Auto Scaling doesn't launch or terminate additional

instances before the previous scaling activity takes effect. After the Auto Scaling group dynamically scales using a simple scaling policy. Auto Scaling waits for the cooldown period to complete before resuming scaling activities. When you manually scale your Auto Scaling group, the default is not to wait for the cooldown period,

but you can override the default and honor the cooldown period. If an instance becomes unhealthy.

Auto Scaling does not wait for the cooldown period to complete before replacing the unhealthy instance

For more information on the Cool down period, please refer to the below URL:

- <http://docs.ws.amazon.com/autoscaling/latest/userguide/Cooldown.html>

NEW QUESTION 245

Your finance supervisor has set a budget of 2000 USD for the resources in AWS. Which of the following is the simplest way to ensure that you know when this threshold is being reached.

- A. Use Cloudwatch events to notify you when you reach the threshold value
- B. Use the Cloudwatch billing alarm to to notify you when you reach the threshold value
- C. Use Cloudwatch logs to notify you when you reach the threshold value
- D. Use SQS queues to notify you when you reach the threshold value

Answer: B

Explanation:

The AWS documentation mentions

You can monitor your AWS costs by using Cloud Watch. With Cloud Watch, you can create billing alerts that notify you when your usage of your services exceeds thresholds that you define. You specify these threshold amounts when you create the billing alerts.

When your usage exceeds these amounts, AWS sends you an

email notification. You can also sign up to receive notifications when AWS prices change. For more information on billing alarms, please refer to the below URL:

- <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/monitor-charges.html>

NEW QUESTION 247

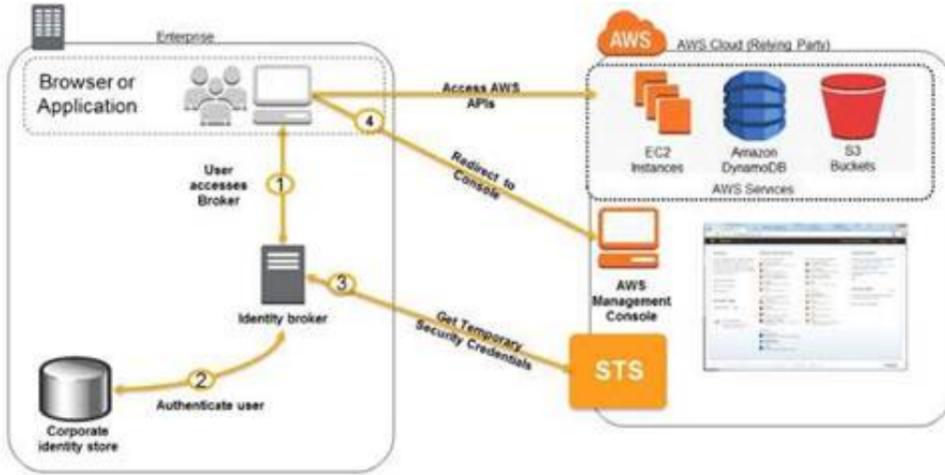
Which of the following will you need to consider so you can set up a solution that incorporates single sign-on from your corporate AD or LDAP directory and restricts access for each user to a designated user folder in a bucket? Choose 3 Answers from the options below

- A. Setting up a federation proxy or identity provider
- B. Using AWS Security Token Service to generate temporary tokens
- C. Tagging each folder in the bucket
- D. Configuring IAM role
- E. Setting up a matching IAM user for every user in your corporate directory that needs access to a folder in the bucket

Answer: ABD

Explanation:

The below diagram showcases how authentication is carried out when having an identity broker. This is an example of a SAML connection, but the same concept holds true for getting access to an AWS resource.



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

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_create_for-idp_saml.html?icmpid=docs_iam_console

<https://aws.amazon.com/blogs/security/writing-iam-policies-grant-access-to-user-specific-folders-in-an-amazon-s3-bucket/>

NEW QUESTION 248

.....

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 First Try
- * DOP-C01 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

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