

Exam Questions DOP-C01

AWS Certified DevOps Engineer- Professional

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



NEW QUESTION 1

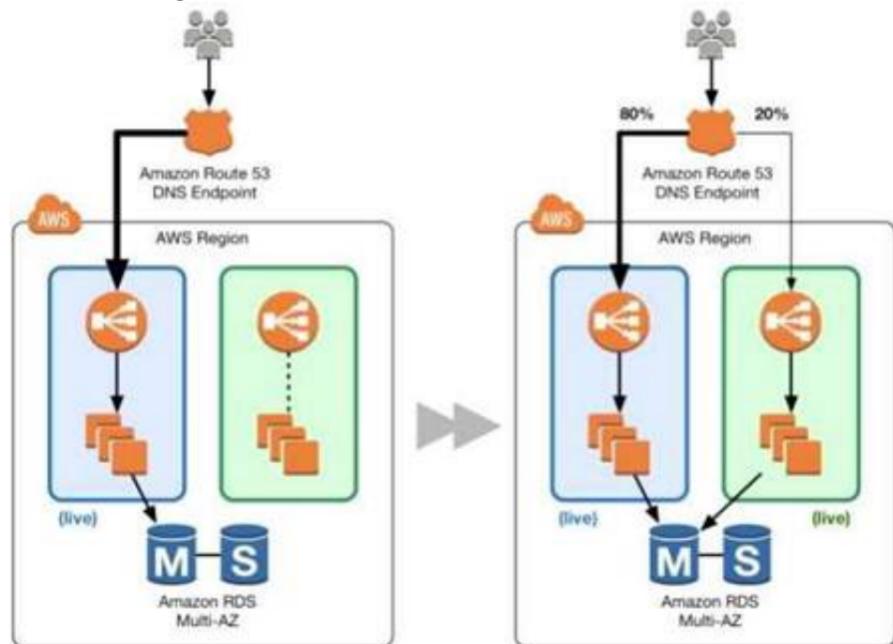
Your application is currently running on Amazon EC2 instances behind a load balancer. Your management has decided to use a Blue/Green deployment strategy. How should you implement this for each deployment?

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

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

Answer: D

Explanation:

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

NEW QUESTION 3

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

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

Answer: CD

Explanation:

These steps are given in the AWS documentation

To migrate data between encrypted and unencrypted volumes

1) Create your destination volume (encrypted or unencrypted, depending on your need).

2) Attach the destination volume to the instance that hosts the data to migrate.

3) Make the destination volume available by following the procedures in Making an Amazon EBS Volume Available for Use. For Linux instances, you can create a mount point at /mnt/destination and mount the destination volume there.

4) Copy the data from your source directory to the destination volume. It may be most convenient to use a bulk-copy utility for this.

To encrypt a volume's data by means of snapshot copying

1) Create a snapshot of your unencrypted CBS volume. This snapshot is also unencrypted.

2) Copy the snapshot while applying encryption parameters. The resulting target snapshot is encrypted.

3) Restore the encrypted snapshot to a new volume, which is also encrypted.

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

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

NEW QUESTION 4

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

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

Answer: D

Explanation:

Option D fulfills all the requirements

1) First is to create a SNS topic for each group so that the required members get the email addresses.

2) Ensure the application uses the HTTPS endpoint and the SDK to publish messages Option A is invalid because the SQS service is not required.

Option B and C are incorrect. As per the requirement we need to provide notification to only those on-call teams who are working in that particular shift when an alarm is triggered. It need not have to be sent to all the on-call teams of the company. With Option B & C, since we are not configuring the SNS topic for each on-call team the notifications will be sent to all the on-call teams. Hence these 2 options are invalid. For more information on setting up notifications, please refer to the below document link: from AWS http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US_SetupSNS.html

NEW QUESTION 5

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

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

Answer: C

Explanation:

Auto Scaling helps you ensure that you have the correct number of Amazon EC2 instances available to handle the load for your application. You create collections of

EC2 instances, called Auto Scaling groups. You can specify the minimum number of instances in each Auto Scaling group, and Auto Scaling ensures that your group

never goes below this size. You can specify the maximum number of instances in each Auto Scaling group, and Auto Scaling ensures that your group never goes above this size. If you specify the desired capacity, either when you create the group or at any time thereafter. Auto Scaling ensures that your group has this many instances. If you specify scaling policies, then Auto Scaling can launch or terminate instances as demand on your application increases or decreases.

Option A and B are invalid because Autoscaling is required to solve the issue to ensure the application can handle high traffic loads.

Option D is invalid because there is no Autoscaling template.

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

NEW QUESTION 6

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 7

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 8

The operations team and the development team want a single place to view both operating system and application logs. How should you implement this using AWS services? Choose two from the options below

- A. Using AWS CloudFormation, create a Cloud Watch Logs LogGroup and send the operating system and application logs of interest using the Cloud Watch Logs Agent.
- B. Using AWS CloudFormation and configuration management, set up remote logging to send events via UDP packets to CloudTrail.
- C. Using configuration management, set up remote logging to send events to Amazon Kinesis and insert these into Amazon CloudSearch or Amazon Redshift, depending on available analytic tools.
- D. Using AWS CloudFormation, merge the application logs with the operating system logs, and use IAM Roles to allow both teams to have access to view console output from Amazon EC2.

Answer: AC

Explanation:

Option B is invalid because Cloudtrail is not designed specifically to take in UDP packets

Option D is invalid because there are already Cloudwatch logs available, so there is no need to have specific logs designed for this.

You can use Amazon CloudWatch Logs to monitor, store, and access your log files from Amazon Elastic Compute Cloud (Amazon EC2) instances, AWS CloudTrail,

and other sources. You can then retrieve the associated log data from CloudWatch Logs. For more information on Cloudwatch logs please refer to the below link:

<http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html> You can use Kinesis to process those logs

For more information on Amazon Kinesis please refer to the below link: <http://docs.aws.amazon.com/streams/latest/dev/introduction.html>

NEW QUESTION 9

You have the following application to be setup in AWS

- 1) A web tier hosted on EC2 Instances
- 2) Session data to be written to DynamoDB
- 3) Log files to be written to Microsoft SQL Server

How can you allow an application to write data to a DynamoDB table?

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

Answer: D

Explanation:

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

NEW QUESTION 10

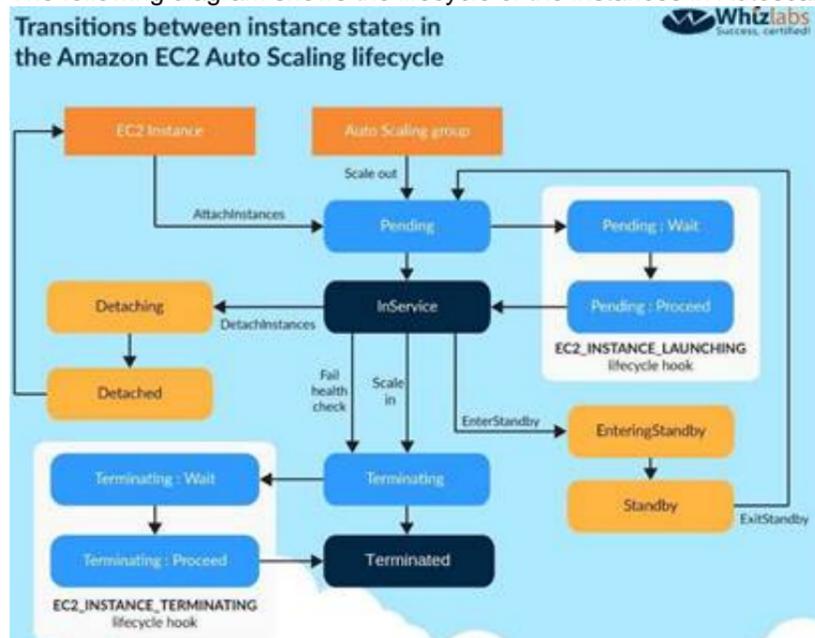
You have an Auto Scaling group of Instances that processes messages from an Amazon Simple Queue Service (SQS) queue. The group scales on the size of the queue. Processing Involves calling a third- party web service. The web service is complaining about the number of failed and repeated calls it is receiving 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 software
- 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 Scaling group, 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.htm> | 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 10

You are administering a continuous integration application that polls version control for changes and then launches new Amazon EC2 instances for a full suite of build tests. What should you do to ensure the lowest overall cost while being able to run as many tests in parallel as possible?

- A. Perform syntax checking on the continuous integration system before launching a new Amazon EC2 instance for build test, unit and integration tests.
- B. Perform syntax and build tests on the continuous integration system before launching the new Amazon EC2 instance unit and integration test
- C. Perform all tests on the continuous integration system, using AWS OpsWorks for unit, integration, and build tests.
- D. Perform syntax checking on the continuous integration system before launching a new AWS Data Pipeline for coordinating the output of unit, integration, and build tests.

Answer: B

Explanation:

Continuous Integration (CI) is a development practice that requires developers to integrate code into a shared repository several times a day. Each check-in is then verified by an automated build, allowing teams to detect problems early.

Option A and D are invalid because you can do build tests on a CI system and not only Syntax tests. And Syntax tests are normally done during coding time and not during the build time.

Option C is invalid because Opswork is ideally not used for build and integration tests.

For an example of a Continuous integration system, please refer to the Jenkins system via the url below

- <https://jenkins.io/>

NEW QUESTION 14

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 scrip
- 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 18

You work for a company that has multiple applications which are very different and built on different programming languages. How can you deploy applications as quickly as possible?

- A. Develop each app in one Docker container and deploy using ElasticBeanstalk
- B. Create a Lambda function deployment package consisting of code and any dependencies
- C. Develop each app in a separate Docker container and deploy using Elastic Beanstalk V
- D. Develop each app in a separate Docker containers and deploy using CloudFormation

Answer: C

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.

Option A is an efficient way to use Docker. The entire idea of Docker is that you have a separate environment for various applications.

Option B is ideally used to running code and not packaging the applications and dependencies Option D is not ideal deploying Docker containers using Cloudformation

For more information on Docker and Clastic Beanstalk, please visit the below URL:

? http://docs.aws.a mazon.com/elasticbeanstalk/latest/dg/create_deploy_docker.html

NEW QUESTION 22

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

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

Answer: BDE

Explanation:

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

Cloud Watch Logs uses your log data for monitoring; so, no code changes are required. For example, you can monitor application logs for specific literal terms (such as "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^ws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html>

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

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

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

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

NEW QUESTION 26

After a daily scrum with your development teams, you've agreed that using Blue/Green style deployments would benefit the team. Which technique should you use

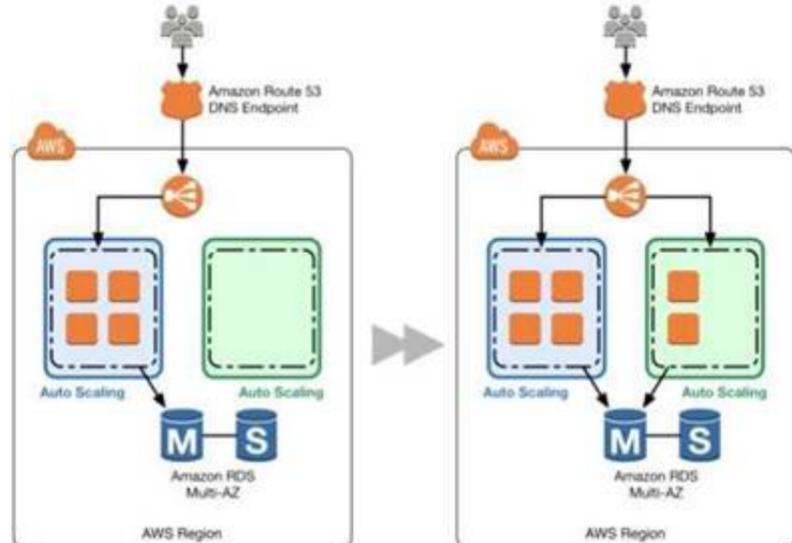
to deliver this new requirement?

- A. Re-deploy your application on AWS Elastic Beanstalk, and take advantage of Elastic Beanstalk deployment types.
- B. Using an AWS CloudFormation template, re-deploy your application behind a load balancer, launch a new AWS CloudFormation stack during each deployment, update your load balancer to send half your traffic to the new stack while you test, after verification update the load balancer to send 100% of traffic to the new stack, and then terminate the old stack.
- C. Create a new AutoScaling group with the new launch configuration and desired capacity same as that of the initial AutoScaling group and associate it with the same load balancer
- D. Once the new AutoScaling group's instances got registered with ELB, modify the desired capacity of the initial AutoScaling group to zero and gradually delete the old AutoScaling group
- E. •>/
- F. Using an AWS OpsWorks stack, re-deploy your application behind an Elastic Load Balancing load balancer and take advantage of OpsWorks stack versioning, during deployment create a new version of your application, tell OpsWorks to launch the new version behind your load balancer, and when the new version is launched, terminate the old OpsWorks stack.

Answer: C

Explanation:

This is given as a practice in the Green Deployment Guides



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 31

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 CLI instances to get those instances which need to be removed from the configuration management system.

NEW QUESTION 32

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

- A. Install an Amazon Cloudwatch Logs Agent on every web server during the bootstrap process
- B. Create a CloudWatch log group and define Metric Filters to create custom metrics that track unique visitors from the streaming web server log
- C. Create a scheduled task on an Amazon EC2 instance that runs every hour to generate a new report based on the Cloudwatch custom metric

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

Answer: AC

Explanation:

You can use the Cloud Watch Logs agent installer on an existing CC2 instance to install and configure the Cloud Watch Logs agent.

For more information, please visit the below link:

- <http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/QuickStartCC2Instance.html>

You can publish your own metrics to Cloud Watch using the AWS CLI or an API. For more information, please visit the below link:

- <http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.html> Amazon Redshift is a fast, fully managed data warehouse that makes it simple and cost-effective to analyze all your data using standard SQL and your existing Business Intelligence (BI) tools. It allows you to run complex analytic queries against petabytes of structured data, using sophisticated query optimization, columnar storage on high-performance local disks, and massively parallel query execution. Most results come back in seconds. For more information on copying data from S3 to redshift, please refer to the below link:
- <http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/dp-copydata-redshift.html>

NEW QUESTION 33

Your application uses Cloud Formation 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 instanc
- C. Only allow UpdateStack permission to 1AM 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 36

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

- A. You need to create an instance profile and associate it with that specific role.
- B. You are not able to associate an 1AM role with an instanceC You won't be able to use that role with an instance unless you also create a user and associate it with that specific role
- C. You won't be able to use that role with an instance unless you also create a usergroup and associate it with that specific role.

Answer: A

Explanation:

An instance profile is a container for an 1AM role that you can use to pass role information to an EC2 instance when the instance starts.

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

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

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

NEW QUESTION 40

You have an application running on Amazon EC2 in an Auto Scaling group. Instances are being bootstrapped dynamically, and the bootstrapping takes over 15 minutes to complete. You find that instances are reported by Auto Scaling as being In Service before bootstrapping has completed. You are receiving application alarms related to new instances before they have completed bootstrapping, which is causing confusion. You find the cause: your application monitoring tool is polling the Auto Scaling Service API for instances that are In Service, and creating alarms for new previously unknown instances. Which of the following will ensure that new instances are not added to your application monitoring tool before bootstrapping is completed?

- A. Create an Auto Scaling group lifecycle hook to hold the instance in a pending: wait state until your bootstrapping is complet
- B. Once bootstrapping is complete, notify Auto Scaling to complete the lifecycle hook and move the instance into a pending:proceed state.
- C. Use the default Amazon Cloud Watch application metrics to monitor your application's 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 stat
- F. Change your application monitoring tool to look for this tag before adding new instances, and the use the Amazon API to set the instance state to 'pending' until bootstrapping is complete.
- G. Increase the desired number of instances in your Auto Scaling group configuration to reduce the time it takes to bootstrap future instances.

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 45

You work for an insurance company and are responsible for the day-to-day operations of your company's online quote system used to provide insurance quotes to members of the public. Your company wants to use the application logs generated by the system to better understand customer behavior. Industry, regulations also require that you retain all application logs for the system indefinitely in order to investigate fraudulent claims in the future. You have been tasked with designing a log management system with the following requirements:

- All log entries must be retained by the system, even during unplanned instance failure.
- The customer insight team requires immediate access to the logs from the past seven days.
- The fraud investigation team requires access to all historic logs, but will wait up to 24 hours before these logs are available.

How would you meet these requirements in a cost-effective manner? Choose three answers from the options below

- A. Configure your application to write logs to the instance's ephemeral disk, because this storage is free and has good write performance.
- B. Create a script that moves the logs from the instance to Amazon S3 once an hour.
- C. Write a script that is configured to be executed when the instance is stopped or terminated and that will upload any remaining logs on the instance to Amazon S3.
- D. Create an Amazon S3 lifecycle configuration to move log files from Amazon S3 to Amazon Glacier after seven days.
- E. Configure your application to write logs to the instance's default Amazon EBS boot volume, because this storage already exist
- F. Create a script that moves the logs from the instance to Amazon S3 once an hour.
- G. Configure your application to write logs to a separate Amazon EBS volume with the "delete on termination" field set to false
- H. Create a script that moves the logs from the instance to Amazon S3 once an hour.
- I. Create a housekeeping script that runs on a T2 micro instance managed by an Auto Scaling group for high availability
- J. The script uses the AWS API to identify any unattached Amazon EBS volumes containing log file
- K. Your housekeeping script will mount the Amazon EBS volume, upload all logs to Amazon S3, and then delete the volume.

Answer: CEF

Explanation:

Since all logs need to be stored indefinitely. Glacier is the best option for this. One can use Lifecycle events to stream the data from S3 to Glacier. Lifecycle configuration enables you to specify the lifecycle management of objects in a bucket. The configuration is a set of one or more rules, where each rule defines an action for Amazon S3 to apply to a group of objects. These actions can be classified as follows:

- Transition actions - In which you define when objects transition to another storage class. For example, you may choose to transition objects to the STANDARD_IA (infrequent access) storage class 30 days after creation, or archive objects to the GLACIER storage class one year after creation.
- Expiration actions - In which you specify when the objects expire. Then Amazon S3 deletes the expired objects on your behalf. For more information on Lifecycle events, please refer to the below link:
- <http://docs.aws.amazon.com/AmazonS3/latest/dev/object-lifecycle-mgmt.htm> | You can use scripts to put the logs onto a new volume and then transfer those logs to S3.

Note:

Moving the logs from CBS volume to S3 we have some custom scripts running in the background. In order to ensure the minimum memory requirements for the OS and the applications for the script to execute we can use a cost effective ec2 instance.

Considering the computing resource requirements of the instance and the cost factor a t2micro instance can be used in this case.

The following link provides more information on various t2 instances. <https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/t2-instances.html>

Question is "How would you meet these requirements in a cost-effective manner? Choose three answers from the options below"

So here user has to choose the 3 options so that the requirement is fulfilled. So in the given 6 options, options C, C and F fulfill the requirement.

"The CC2s use CBS volumes and the logs are stored on CBS volumes those are marked for non-termination" - is one of the ways to fulfill requirement. So this shouldn't be an issue.

NEW QUESTION 46

You need to implement Blue/Green Deployment for several multi-tier web applications. Each of them has its individual infrastructure:

Amazon Elastic Compute Cloud (EC2) front-end servers, Amazon ElastiCache clusters, Amazon Simple Queue Service (SQS) queues, and Amazon Relational Database (RDS) Instances.

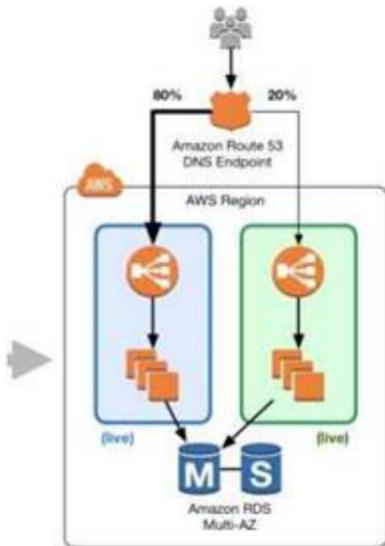
Which combination of services would give you the ability to control traffic between different deployed versions of your application?

- A. Create one AWS Elastic Beanstalk application and all AWS resources (using configuration files inside the application source bundle) for each web application
- B. New versions would be deployed using Elastic Beanstalk environments and using the Swap URLs feature.
- C. Using AWS CloudFormation templates, create one Elastic Beanstalk application and all AWS resources (in the same template) for each web application
- D. New versions would be deployed using AWS CloudFormation templates to create new Elastic Beanstalk environments, and traffic would be balanced between them using weighted Round Robin (WRR) records in Amazon Route 53. >/
- E. Using AWS CloudFormation templates, create one Elastic Beanstalk application and all AWS resources (in the same template) for each web application
- F. New versions would be deployed updating a parameter on the CloudFormation template and passing it to the cfn-hup helper daemon, and traffic would be balanced between them using Weighted Round Robin (WRR) records in Amazon Route 53.
- G. Create one Elastic Beanstalk application and all AWS resources (using configuration files inside the application source bundle) for each web application
- H. New versions would be deployed updating the Elastic Beanstalk application version for the current Elastic Beanstalk environment.

Answer: B

Explanation:

This is an example of Blue green deployment



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

When it's time to promote the green environment/stack into production, update DNS records to point to the green environment/stack's load balancer. You can also do this DNS flip gradually by using the Amazon Route 53 weighted routing policy. For more information on Blue green deployment, please refer to the link:

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

NEW QUESTION 47

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.

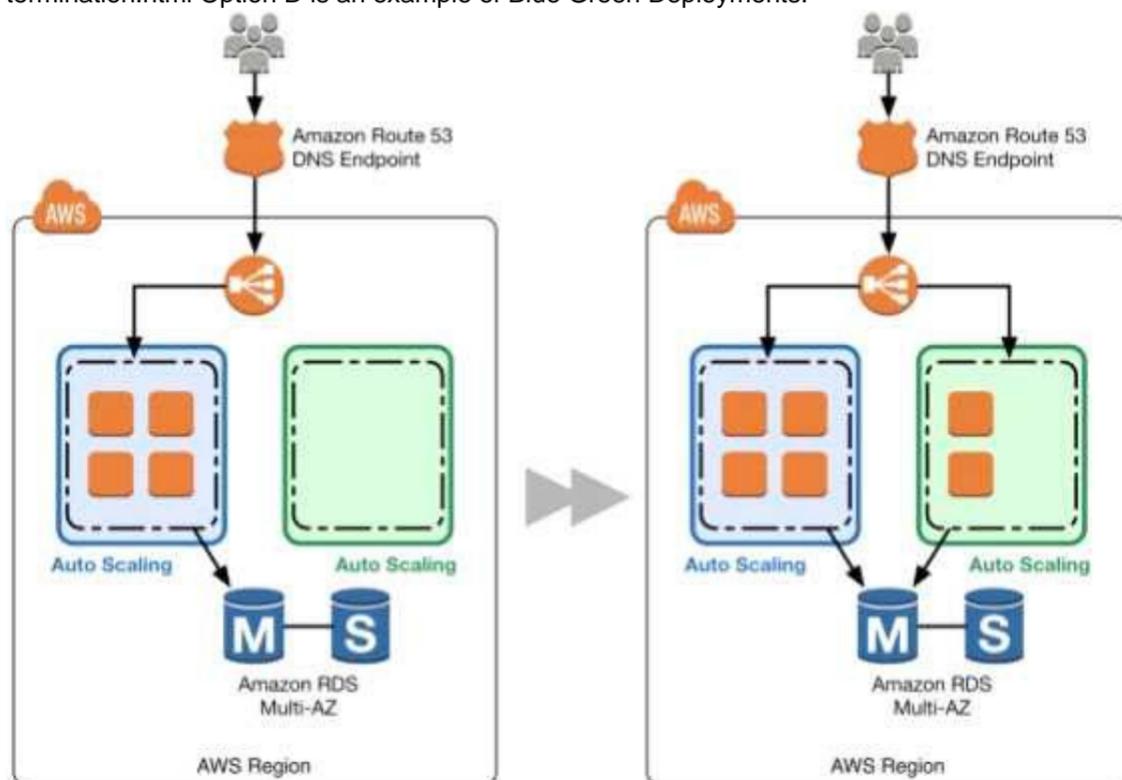
- A. Use Newest Instance to phase out all instances that use the previous configuration.
- B. Attach an additional ELB to your Auto Scaling configuration and phase in newer instances while removing older instances.
- C. Use OldestLaunchConfiguration to phase out all instances that use the previous configuration.
- 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 49

You are designing a system which needs, at a minimum, 8 m4.large instances operating to service traffic. When designing a system for high availability in the us-east-1 region, which has 6 Availability Zones, your company needs to be able to handle the death of a full availability zone. How should you distribute the servers, to save as much cost as possible, assuming all of the EC2 nodes are properly linked to an ELB? Your VPC account can utilize us-east-1's AZ's a through f, inclusive.

- A. 3 servers in each of AZ's a through d, inclusive
- B. 8 servers in each of AZ's a and b.
- C. 2 servers in each of AZ's a through e, inclusive.
- D. 4 servers in each of AZ's a through conclusive.

Answer: C

Explanation:

The best way is to distribute the instances across multiple AZ's to get the best and avoid a disaster scenario. With this scenario, you will always a minimum of more than 8 servers even if one AZ were to go down. Even though A and D are also valid options, the best option when it comes to distribution is Option C. For more information on High Availability and Fault tolerance, please refer to the below link:
https://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_ftha_04.pdf

NEW QUESTION 54

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

- A. Use Direct Connect to route traffic to the on-premise locatio
- B. In DirectConnect, configure the amount of traffic to be routed to the on-premise location.
- C. Implement a Route 53 failover routing policy that sends traffic back to the on-premises application if the AWS application fails.
- D. Configure an Elastic Load Balancer to distribute the traffic between the on-premises application and the AWS application.
- E. Implement a Route 53 weighted routing policy that distributes the traffic between your on- premises application and the AWS application depending on weight.

Answer: D

Explanation:

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

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

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

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

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

NEW QUESTION 57

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

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

Answer: B

Explanation:

AWS OpsWorks is a configuration management service that uses Chef, an automation platform that treats server configurations as code. OpsWorks uses Chef to automate how servers are configured, deployed, and managed across your Amazon Clastic Compute Cloud (Amazon CC2) instances or on-premises compute environments. OpsWorks has two offerings, AWS Opsworks for Chef Automate, and AWS OpsWorks Stacks.

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

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

NEW QUESTION 58

You have an application running on an Amazon EC2 instance and you are using 1AM roles to securely access AWS Service APIs. How can you configure your application running on that instance to retrieve the API keys for use with the AWS SDKs?

- A. When assigning an EC2IAM role to your instance in the console, in the "Chosen SDK" drop-down list, select the SDK that you are using, and the instance will configure the correct SDK on launch with the API keys.
- B. Within your application code, make a GET request to the 1AM Service API to retrieve credentials for your user.
- C. When using AWS SDKs and Amazon EC2 roles, you do not have to explicitly retrieve API keys, because the SDK handles retrieving them from the Amazon EC2 MetaData service.
- D. Within your application code, configure the AWS SDK to get the API keys from environment variables, because assigning an Amazon EC2 role stores keys in environment variables on launch.

Answer: C

Explanation:

IAM roles are designed so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that

the applications use. Instead of creating and distributing your AWS credentials, you can delegate permission to make API requests using 1AM roles

For more information on Roles for CC2 please refer to the below link: <http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

NEW QUESTION 62

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

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

Answer: C

Explanation:

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

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

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

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

Some documents are frequently accessed for a limited period of time. After that, these documents are less frequently accessed. Over time, you might not need real-time access to these objects, but your organization or regulations might require you to archive them for a longer period and then optionally delete them later. You might also upload some types of data to Amazon S3 primarily for archival purposes, for example digital media archives, financial and healthcare records, raw genomics sequence data, long-term database backups, and data that must be retained for regulatory compliance.

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

Note:

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

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

NEW QUESTION 67

You are planning on using the Amazon RDS facility for Fault tolerance for your application. How does Amazon RDS Multi Availability Zone model work

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

Answer: A

Explanation:

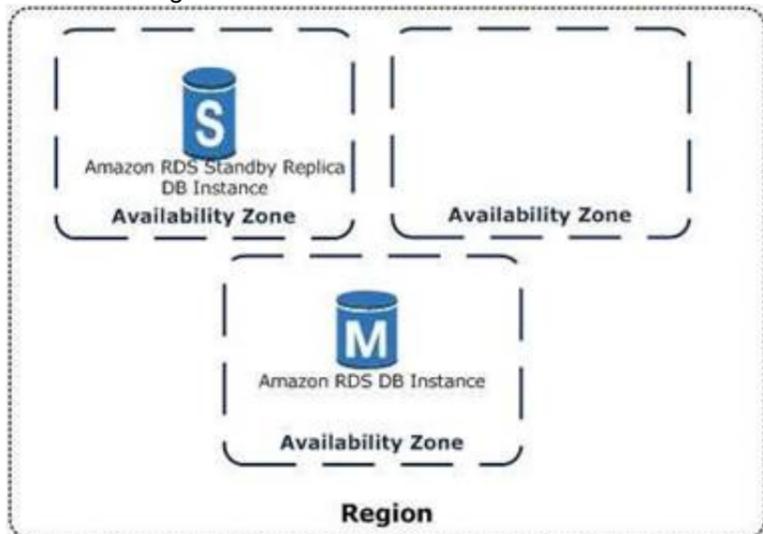
Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB)

Instances, making them a natural fit for production database

workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable.

In case of an infrastructure failure, Amazon RDS performs an automatic failover to the standby (or to a read replica in the case of Amazon Aurora), so that you can resume database operations as soon as the failover is complete.

The below diagram from the AWS documentation shows how this is configured



Option B is invalid because the replication is synchronous.

Option C and D are invalid because this is built around AZ and not regions. For more information on Multi-AZ RDS, please visit the below URL:

<https://aws.amazon.com/rds/details/multi-az/>

NEW QUESTION 69

You need to perform ad-hoc analysis on log data, including searching quickly for specific error codes and reference numbers. Which should you evaluate first?

- A. AWS Elasticsearch Service
- B. AWSRedShift
- C. AWSEMR
- D. AWSDynamoDB

Answer: A

Explanation:

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

Amazon

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

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

NEW QUESTION 71

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

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

Answer: A

Explanation:

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

Reserved Instances are not physical instances, but rather a

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

For more information, please refer to the below link:

- <https://aws.amazon.com/about-aws/whats-new/2011/12/01/New-Amazon-CC2-Reserved-Instances-Options-Now-Available/>
- <https://aws.amazon.com/blogs/aws/reserved-instance-options-for-amazon-ec2/>
- <http://docs.aws.amazon.com/AWSEC2/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 72

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 76

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 77

You are creating an application which stores extremely sensitive financial information. All information in the system must be encrypted at rest and in transit. Which of these is a violation of this policy?

- A. ELB SSL termination.
- B. ELB Using Proxy Protocol v1.
- C. CloudFront Viewer Protocol Policy set to HTTPS redirection.

D. Telling S3 to use AES256 on the server-side.

Answer: A

Explanation:

If you use SSL termination, your servers will always get non-secure connections and will never know whether users used a more secure channel or not. If you are using Elastic beanstalk to configure the ELB, you can use the below article to ensure end to end encryption.
<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/configuring-https-endtoend.html>

NEW QUESTION 80

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

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

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

You can generate and download a credential report that lists all users in your account and the status of their various credentials, including passwords, access keys, and MFA devices. You can get a credential report from the AWS Management Console, the AWS SDKs and Command Line Tools, or the 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 81

Your application consists of 10% writes and 90% reads. You currently service all requests through a Route53 Alias Record directed towards an AWS ELB, which sits in front of an EC2 Auto Scaling Group. Your system is getting very expensive when there are large traffic spikes during certain news events, during which many more people request to read similar data all at the same time. What is the simplest and cheapest way to reduce costs and scale with spikes like this?

- A. Create an S3 bucket and asynchronously replicate common requests responses into S3 object
- B. When a request comes in for a precomputed response, redirect to AWS S3.
- C. Create another ELB and Auto Scaling Group layer mounted on top of the other system, adding a tier to the system
- D. Serve most read requests out of the top layer.
- E. Create a CloudFront Distribution and direct Route53 to the Distribution
- F. Use the ELB as an Origin and specify Cache Behaviours to proxy cache requests which can be served late.
- G. Create a Memcached cluster in AWS ElastiCache
- H. Create cache logic to serve requests which can be served late from the in-memory cache for increased performance.

Answer: C

Explanation:

Use CloudFront distribution for distributing the heavy reads for your application. You can create a zone apex record to point to the Cloudfront distribution.

You can control how long your objects stay in a CloudFront cache before CloudFront forwards another request to your origin. Reducing the duration allows you to serve dynamic content. Increasing the duration means your users get better performance because your objects are more likely to be served directly from the edge cache. A longer duration also reduces the load on your origin.

For more information on Cloudfront object expiration, please visit the below URL:

- <http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Expiration.html>

NEW QUESTION 84

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

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

Answer: A

Explanation:

You can use Lambda functions as triggers for your Amazon DynamoDB table. Triggers are custom actions you take in response to updates made to the

DynamoDB table. To create a trigger, first you enable Amazon DynamoDB Streams for your table. Then, you write a Lambda function to process the updates published to the stream.

For more information on DynamoDB with Lambda, please visit the below URL: <http://docs.aws.amazon.com/lambda/latest/dg/with-ddb.html>

NEW QUESTION 85

If I want CloudFormation 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 CloudFormation 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 87

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 bplex 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 88

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 90

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

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

Answer: B

Explanation:

Answer - B

Some of the best practices for Cloudformation are

- Created Nested stacks

As your infrastructure grows, common patterns can emerge in which you declare the same components in each of your templates. You can separate out these common components and create dedicated templates for them. That way, you can mix and match different templates but use nested stacks to create a single,

unified stack. Nested stacks are stacks that create other stacks. To create nested stacks, use the `AWS::CloudFormation::StackResource` in your template to reference other templates.

- Reuse Templates

After you have your stacks and resources set up, you can reuse your templates to replicate your infrastructure in multiple environments. For example, you can create environments for development, testing, and production so that you can test changes before implementing them into production. To make templates reusable, use the parameters, mappings, and conditions sections so that you can customize your stacks when you create them. For example, for your development environments, you can specify a lower-cost instance type compared to your production environment, but all other configurations and settings remain the same. For more information on CloudFormation best practices, please visit the below URL:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html>

NEW QUESTION 91

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 95

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 96

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

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

Answer: D

Explanation:

Failover routing lets you route traffic to a resource when the resource is healthy or to a different resource when the first resource is unhealthy. The primary and secondary resource record sets can route traffic to anything from an Amazon S3 bucket that is configured as a website to a complex tree of records.

For more information on Route53 Failover Routing, please visit the below URL:

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

NEW QUESTION 100

You need your CI to build AMIs with code pre-installed on the images on every new code push. You need to do this as cheaply as possible. How do you do this?

- A. Bid on spot instances just above the asking price as soon as new commits come in, perform all instance configuration and setup, then create an AMI based on the spot instance.
- B. Have the CI launch a new on-demand EC2 instance when new commits come in, perform all instance configuration and setup, then create an AMI based on the on-demand instance.
- C. Purchase a Light Utilization Reserved Instance to save money on the continuous integration machine.
- D. Use these credits whenever you create AMIs on instances.
- E. When the CI instance receives commits, attach a new EBS volume to the CI machine.
- F. Perform all setup on this EBS volume so you don't need

Answer: A

Explanation:

Amazon EC2 Spot instances allow you to bid on spare Amazon EC2 computing capacity. Since Spot instances are often available at a discount compared to On-Demand pricing, you can significantly reduce the cost of running your applications, grow your application's compute capacity and throughput for the same budget, and enable new types of cloud computing applications.

For more information on Spot Instances, please visit the below URL: <https://aws.amazon.com/ec2/spot/>

NEW QUESTION 105

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.
- F. Delete the IAM Role associated with the keys after every 2 months. 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 107

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 110

You currently have EC2 Instances hosting an application. These instances are part of an Auto Scaling Group. You now want to change the instance type of the EC2 Instances. How can you manage the deployment with the least amount of downtime?

- A. Terminate the existing Auto Scaling group
- B. Create a new launch configuration with the new Instance type
- C. Attach that to the new Auto Scaling Group.
- D. Use the `AutoScalingRollingUpdate` policy on CloudFormation Template Auto Scaling group
- E. Use the Rolling Update feature which is available for EC2 Instances.
- F. Manually terminate the instances, launch new instances with the new instance type and attach them to the Auto Scaling group

Answer: B

Explanation:

The `AWS::AutoScaling::AutoScalingGroup` resource supports an `UpdatePolicy` attribute. This is used to define how an Auto Scaling group resource is updated when

an update to the Cloud Formation stack occurs. A common approach to updating an Auto Scaling group is to perform a rolling update, which is done by specifying the AutoScalingRollingUpdate policy. This retains the same Auto Scaling group and replaces old instances with new ones, according to the parameters specified. For more information on AutoScaling Rolling Update, please refer to the below link:
 • <https://aws.amazon.com/premiumsupport/knowledge-center/auto-scaling-group-rolling-updates/>

NEW QUESTION 114

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- Nodejs, IIS, etc
- C. Choosing the type of Notification - SNS or SQS
- D. Choosing whether you want a highly available environment or not

Answer: ABD

Explanation:

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



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

NEW QUESTION 118

You are deciding on a deployment mechanism for your application. Which of the following deployment mechanisms provides the fastest rollback after failure.

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

Answer: D

Explanation:

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

NEW QUESTION 122

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

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

Answer: C

Explanation:

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

NEW QUESTION 126

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. S3Lifecycle
- D. Route53

Answer: B

Explanation:

The AWS documentation mentions

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

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

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

NEW QUESTION 131

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 136

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

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

Answer: ABD

Explanation:

The AWS documentation mentions

I AM supports AWS CodeCommit with three types of credentials:

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

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

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

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

NEW QUESTION 139

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

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

Answer: B

Explanation:

The AWS Documentation mentions

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

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

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

NEW QUESTION 143

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

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

Answer: B

Explanation:

The AWS Documentation mentions

AWS OpsWorks Stacks provides built-in support for Memcached. However, if Redis better suits your requirements, you can customize your stack so that your application servers use OastlCache Redis. Although it works with Redis clusters, AWS clearly specifies that AWS Opsworks stacks provide built in support for Memcached.

Amazon OastlCache is an AWS service that makes it easy to provide caching support for your application server, using either the Memcached or Redis caching engines. OastlCache can be used to improve the application server performance running on AWS Opsworks stacks.

For more information on Opswork and Cache engines please refer to the below link:

- <http://docs.aws.amazon.com/opsworks/latest/userguide/other-services-redis.html>

NEW QUESTION 146

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

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

Answer: D

Explanation:

All of the services can be used to implement Devops in your company

- 1) AWS Elastic Beanstalk, an easy-to-use service for deploying and scaling web applications and services developed with Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker on servers such as Apache, Nginx, Passenger, and IIS.
- 2) AWS Ops Works, a configuration management service that helps you configure and operate applications of all shapes and sizes using Chef
- 3) AWS Cloud Formation, which is an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable fashion.

For more information on AWS Devops please refer to the below link:

- <http://docs.aws.amazon.com/devops/latest/gsg/welcome.html>

NEW QUESTION 150

Which of the following service can be used to provision ECS Cluster containing following components in an automated way:

- 1) Application Load Balancer for distributing traffic among various task instances running in EC2 Instances
- 2) Single task instance on each EC2 running as part of auto scaling group
- 3) Ability to support various types of deployment strategies

- A. SAM
- B. Opswork
- C. Elastic beanstalk
- D. CodeCommit

Answer: C

Explanation:

You can create docker environments that support multiple containers per Amazon EC2 instance with multi-container Docker platform for Elastic Beanstalk-Elastic Beanstalk uses Amazon Elastic Container Service (Amazon ECS) to coordinate container deployments to multi-container Docker environments. Amazon ECS provides tools to manage a cluster of instances running Docker containers. Elastic Beanstalk takes care of Amazon ECS tasks including cluster creation, task definition, and execution Please refer to the below AWS documentation: https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_ecs.html

NEW QUESTION 151

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

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

Answer: C

Explanation:

The AWS Documentation specifies

A Dockerrun.aws.json file is an Elastic Beanstalk-specific JSON file that describes how to deploy a set of Docker containers as an Elastic Beanstalk application. You can use a Dockerrun.aws.json file for a multicontainer Docker environment.

Dockerrun.aws.json describes the containers to deploy to each container instance in the environment as well as the data volumes to create on the host instance for the containers to mount.

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

- http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_v2config.html

NEW QUESTION 153

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.

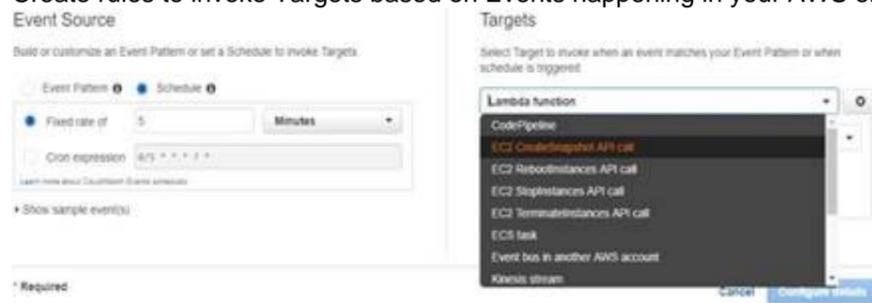
- A. Create a powershell script which uses the AWS CLI to get the volumes and then run the script as a cron job.
- B. Use the AWSConfig service to create a snapshot of the AWS Volumes
- C. Use the AWS CodeDeploy service to create a snapshot of the AWS Volumes
- D. Use Cloudwatch Events to trigger the 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 Cvents delivers a near real-time stream of system events that describe changes in Amazon Web Services (AWS) resources. Using simple rules that you can quickly set up, you can match events and route them to one or more target functions or streams. Cloud Watch Cvents becomes aware of operational changes as they occur. Cloud Watch Cvents responds to these operational changes and takes corrective action as necessary, by sending messages to respond to the environment, activating functions, making changes, and capturing state information. For more information on Cloud watch Cvents, please visit the below U RL: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/events/WhatIsCloudWatchCvents.html>

NEW QUESTION 157

Your company is planning to setup a wordpress application. The wordpress application will connect to a MySQL database. Part of the requirement is to ensure that the database environment is fault tolerant and highly available. Which of the following 2 options individually can help fulfil this requirement.

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

Answer: AD

Explanation:

One way to ensure high availability and fault tolerant environments is to ensure Instances are located across multiple availability zones. Hence if you are hosting MySQL yourself, ensure you have instances spread across multiple AZ's

The AWS Documentation mentions the following about the multi-AZ feature

Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. Amazon RDS uses several different technologies to provide failover support. Multi-AZ deployments for Oracle, PostgreSQL, MySQL, and MariaDB DB instances use Amazon's failover technology

For more information on AWS Multi-AZ deployments, please visit the below URL:

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

NEW QUESTION 159

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

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

Answer: A

Explanation:

The AWS Documentation mentions

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

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

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

NEW QUESTION 162

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

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

Answer: BC

Explanation:

Using User data and cloud-init directives you can always ensure you download the latest version of the LAMP stack and give it to the development teams. With AMI's you will always have the same version and will need to create an AMI everytime the version of the LAMP stack changes. The AWS Documentation mentions When you launch an instance in Amazon 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 165

You are working as an AWS Devops admin for your company. You are in-charge of building the infrastructure for the company's development teams using Cloudformation. The template will include building the VPC and networking components, installing a LAMP stack and securing the created resources. As per the AWS best practices what is the best way to design this template

- A. Create a single cloudformation template to create all the resources since it would be easier from the maintenance perspective.
- B. Create multiple cloudformation templates based on the number of VPC's in the environment.
- C. Create multiple cloudformation templates based on the number of development groups in the environment.
- D. Create multiple cloudformation templates for each set of logical resources, one for networking, the other for LAMP stack creation.

Answer: D

Explanation:

Creating multiple cloudformation templates is an example of using nested stacks. The advantage of using nested stacks is given below as per the AWS documentation As your infrastructure grows, common patterns can emerge in which you declare the same components in each of your templates. You can separate out these common components and create dedicated templates for them. That way, you can mix and match different templates but use nested stacks to create a single, unified stack. Nested stacks are stacks that create other stacks. To create nested stacks, use the `AWS::CloudFormation::Stack` resource in your template to reference other templates. For more information on Cloudformation best practices, please refer to the below link: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html>

NEW QUESTION 170

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 173

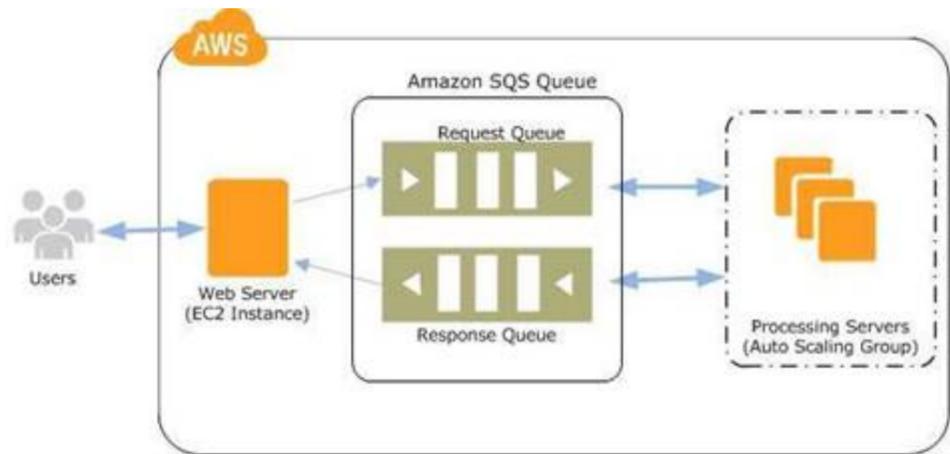
You are having a web and worker role infrastructure defined in AWS using Amazon EC2 resources. You are using SQS to manage the jobs being sent by the web role. Which of the following is the right way to ensure the worker processes are adequately setup to handle the number of jobs sent by the web role

- A. Use Cloudwatch monitoring to check the size of the queue and then scale out SQS to ensure that it can handle the right number of jobs
- B. Use ELB to ensure that the load is evenly distributed to the set of web and worker instances
- C. Use Route53 to ensure that the load is evenly distributed to the set of web and worker instances
- D. Use Cloudwatch monitoring to check the size of the queue and then scale out using Autoscaling to ensure that it can handle the right number of jobs

Answer: D

Explanation:

The below diagram shows how SQS can be used to manage the communication between the Web and worker roles. The number of messages in the SQS queue can be used to determine the number of instances that should be there in the AutoScaling Group.



For more information on SQS and Autoscaling, please refer to the below U RL: <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-using-sqs-queue.html>

NEW QUESTION 178

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

What can you do to ensure the instances are placed in the right state after the bootstrapping process is complete?

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

Answer: B

Explanation:

The AWS Documentation mentions the following

3. If you finish the custom action before the timeout period ends, use the complete-lifecycle-action command so that the Auto Scalinggroup can continue launching

or terminating the instance. You can specify the lifecycle action token, as shown in the following command:

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

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

For more information on lifecycle hooks, please refer to the below URL:

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

NEW QUESTION 183

Which of the following is incorrect when it comes to using the instances in an Opswork stack?

- A. In a stack you can use a mix of both Windowsand Linux operating systems
- B. You can start and stop instances manually in a stack
- C. You can use custom AMI'S as long as they are based on one of the AWS OpsWorks Stacks- supported AMIs
- D. You can use time-based automatic scaling with any stack

Answer: A

Explanation:

The AWS documentation mentions the following about Opswork stack

- A stack's instances can run either Linux or Windows.

A stack can have different Linux versions or distributions on different instances, but you cannot mix Linux and Windows instances.

- You can use custom AMIs (Amazon Machine Images), but they must be based on one of the AWS Ops Works Stacks-supported AMIs
- You can start and stop instances manually or have AWS OpsWorks Stacks automatically scale the number of instances. You can use time-based automatic scaling with any stack; Linux stacks also can use load-based scaling.
- In addition to using AWS OpsWorks Stacks to create Amazon EC2 instances, you can also register instances with a Linux stack that were created outside of AWS OpsWorks Stacks.

For more information on Opswork stacks, please visit the below link: <http://docs.aws.amazon.com/opsworks/latest/userguide/workinginstances-os.html>

NEW QUESTION 188

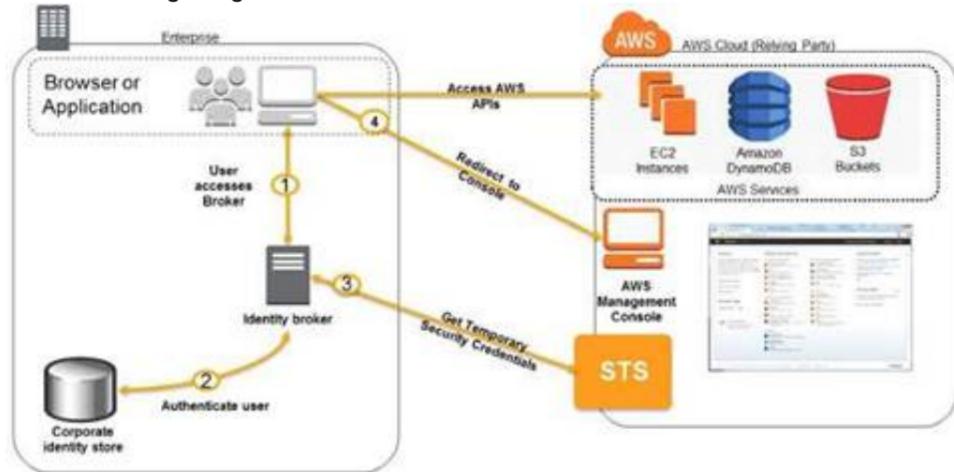
Which of the following will you need to consider so you can set up a solution that incorporates single sign-on from your corporate AD or LDAP directory and restricts access for each user to a designated user folder in a bucket? Choose 3 Answers from the options below

- 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 1AM role
- E. Setting up a matching 1AM 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 190

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

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

Answer: C

Explanation:

In Rolling deployment, you can mention a new set of servers which can replace the existing set of servers. This replacement will happen in a phased out manner. Since there is a requirement to not swap URL's, you must not use Blue Green deployments.

For more information on the differences between Rolling Updates and Blue Green deployments, please refer to the below URL:

- <https://cloudnative.io/docs/blue-green-deployment/>

NEW QUESTION 193

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

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

Answer: D

Explanation:

The AWS Documentation mentions the following on package updates for the Elastic beanstalk environment

You can configure your environment to apply minor and patch version updates automatically during a configurable weekly maintenance window with Managed Platform Updates. Elastic Beanstalk applies managed updates with no downtime or reduction in capacity, and cancels the update immediately if instances running your application on the new version fail health checks.

For more information on Elastic beanstalk managed updates please refer to the URL: <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environment-platform-update-managed.html>

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.platform.upgrade.html>

NEW QUESTION 194

Of the 6 available sections on a Cloud Formation template (Template Description Declaration, Template Format Version Declaration, Parameters, Resources, Mappings, Outputs), which is the only one required for a CloudFormation template to be accepted? Choose an answer from the options below

- A. Parameters
- B. Template Declaration
- C. Mappings
- D. Resources

Answer: D

Explanation:

If you refer to the documentation, you will see that Resources is the only mandatory field

Specifies the stack resources and their properties, such as an Amazon Elastic Compute Cloud instance or an Amazon Simple Storage Service bucket.

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

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

NEW QUESTION 198

You were just hired as a DevOps Engineer for a startup. Your startup uses AWS for 100% of their infrastructure. They currently have no automation at all for deployment, and they have had many failures while trying to deploy to production. The company has told you deployment process risk mitigation is the most important thing now, and you have a lot of budget for tools and AWS resources. Their stack includes a 2-tier API with data stored in DynamoDB or S3, depending on type. The Compute layer is EC2 in Auto Scaling Groups. They use Route53 for DNS pointing to an ELB. An ELB balances load across the EC2 instances. The scaling group properly varies between 4 and 12 EC2 servers. Which of the following approaches, given this company's stack and their priorities, best meets the company's needs?

- A. Model the stack in AWS Elastic Beanstalk as a single Application with multiple Environment
- B. Use Elastic Beanstalk's Rolling Deploy option to progressively roll out application code changes when promoting across environments.
- C. Model the stack in three CloudFormation templates: Data layer, compute layer, and networking layer
- D. Write stack deployment and integration testing automation following Blue-Green methodology
- E. •>/
- F. Model the stack in AWS OpsWorks as a single Stack, with 1 compute layer and its associated EL
- G. Use Chef and App Deployments to automate Rolling Deployment.
- H. Model the stack in 1 CloudFormation template, to ensure consistency and dependency graph resolution
- I. Write deployment and integration testing automation following Rolling Deployment methodologies.

Answer: B

Explanation:

Here you are using 2 of the best practices for deployment, one is Blue Green Deployments and the other is using Nested Cloudformation stacks.

The AWS Documentation mentions the below on nested stacks

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

For more information on Cloudformation best practises, please visit the link:

- <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html> For more information on Blue Green Deployment, please visit the link:
- https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 202

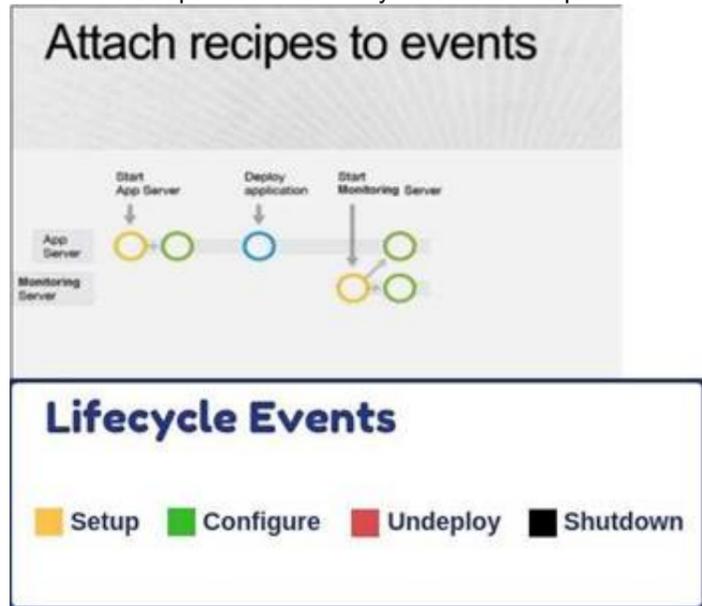
Which of the following are Lifecycle events available in Opswork? Choose 3 answers from the options below

- A. Setup
- B. Decommission
- C. Deploy
- D. Shutdown

Answer: ACD

Explanation:

Below is a snapshot of the Lifecycle events in Opswork.



For more information on Lifecycle events, please refer to the below URL:

- <http://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-events.html>

NEW QUESTION 205

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

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

Answer: B

Explanation:

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

reliable.

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

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

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

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

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

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

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

NEW QUESTION 206

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

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

Answer: B

Explanation:

The AWS documentation mentions the following

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

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

Specify whether messages are always stored in Amazon S3 or only when a message's size exceeds 256 KB.

Send a message that references a single message object stored in an Amazon S3 bucket. Get the corresponding message object from an Amazon S3 bucket.

Delete the corresponding message object from an Amazon S3 bucket. For more information on SQS and sending larger messages please visit the link

NEW QUESTION 210

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

DOP-C01 Practice Exam Features:

- * DOP-C01 Questions and Answers Updated Frequently
- * DOP-C01 Practice Questions Verified by Expert Senior Certified Staff
- * DOP-C01 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * DOP-C01 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year