

Exam Questions DOP-C01

AWS Certified DevOps Engineer- Professional

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



NEW QUESTION 1

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

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

Answer: A

Explanation:

The AWS::AutoScaling::AutoScalingGroup resource supports an UpdatePolicy attribute. This is used to define how an Auto 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 updates, please refer to the below link: <https://aws.amazon.com/premiumsupport/knowledge-center/auto-scaling-group-rolling-updates/>

NEW QUESTION 2

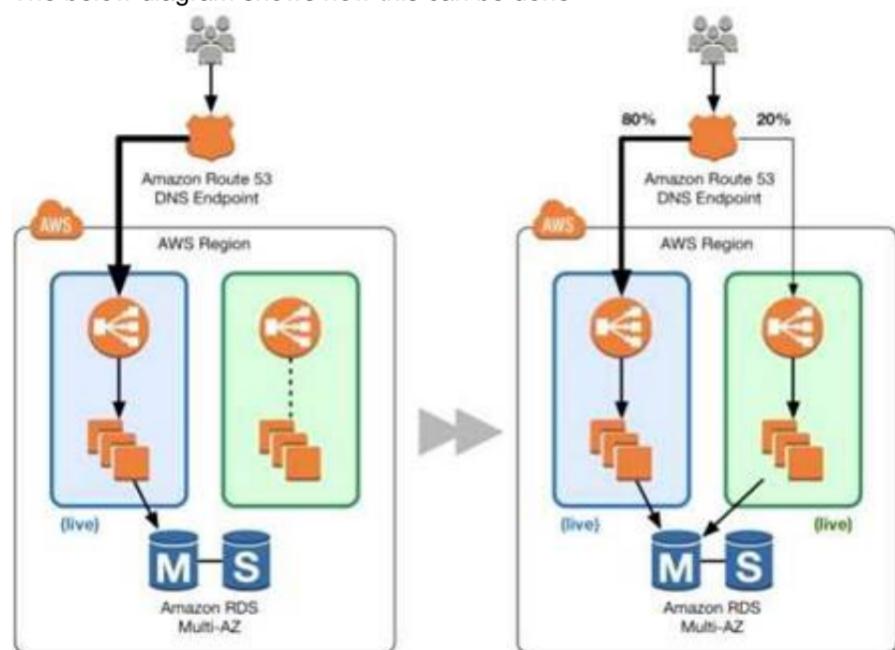
Your application is currently running on Amazon EC2 instances behind a load balancer. Your management has decided to use a Blue/Green deployment strategy. How should you implement this for each deployment?

- 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 3

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 4

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

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

Answer: C

Explanation:

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

NEW QUESTION 5

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 6

You have an application consisting of a stateless web server tier running on Amazon EC2 instances behind load balancer, and are using Amazon RDS with read replicas. Which of the following methods should you use to implement a self-healing and cost-effective architecture? Choose 2 answers from the options given below

- A. Set up a third-party monitoring solution on a cluster of Amazon EC2 instances in order to emit custom Cloud Watch metrics to trigger the termination of unhealthy Amazon EC2 instances.
- B. Set up scripts on each Amazon EC2 instance to frequently send ICMP pings to the load balancer in order to determine which instance is unhealthy and replace it.
- C. Set up an Auto Scaling group for the web server tier along with an Auto Scaling policy that uses the Amazon RDS DB CPU utilization Cloud Watch metric to scale the instances.
- D. Set up an Auto Scaling group for the web server tier along with an Auto Scaling policy that uses the Amazon EC2 CPU utilization CloudWatch metric to scale the instances.
- E. Use a larger Amazon EC2 instance type for the web server tier and a larger DB instance type for the data storage layer to ensure that they don't become unhealthy.
- F. Set up an Auto Scaling group for the database tier along with an Auto Scaling policy that uses the Amazon RDS read replica lag CloudWatch metric to scale out the Amazon RDS read replicas.
- G. Use an Amazon RDS Multi-AZ deployment.

Answer: DG

Explanation:

The scaling of EC2 Instances in the Autoscaling group is normally done with the metric of the CPU utilization of the current instances in the Autoscaling group. For more information on scaling in your Autoscaling Group, please refer to the below link:

• <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scaling-simple-step.html>

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. For more information on RDS Multi-AZ please refer to the below link:

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

Option A is invalid because if you already have in-built metrics from Cloudwatch, why would you want to spend more in using a third-party monitoring solution.

Option B is invalid because health checks are already a feature of AWS CLB

Option C is invalid because the database CPU usage should not be used to scale the web tier.

Option D is invalid because increasing the instance size does not always guarantee that the solution will not become unhealthy.

Option E is invalid because increasing Read-Replica's will not suffice for write operations if the primary DB fails.

NEW QUESTION 7

You have a code repository that uses Amazon S3 as a data store. During a recent audit of your security controls, some concerns were raised about maintaining the integrity of the data in the Amazon S3 bucket. Another concern was raised around securely deploying code from Amazon S3 to applications running on Amazon EC2 in a virtual private cloud. What are some measures that you can implement to mitigate these concerns? Choose two answers from the options given below.

A. Add an Amazon S3 bucket policy with a condition statement to allow access only from Amazon EC2 instances with RFC 1918 IP addresses and enable bucket versioning.

B. Add an Amazon S3 bucket policy with a condition statement that requires multi-factor authentication in order to delete objects and enable bucket versioning.

C. Use a configuration management service to deploy AWS Identity and Access Management user credentials to the Amazon EC2 instance

D. Use these credentials to securely access the Amazon S3 bucket when deploying code.

E. Create an Amazon Identity and Access Management role with authorization to access the Amazon S3 bucket, and launch all of your application's Amazon EC2 instances with this role.

F. Use AWS Data Pipeline to lifecycle the data in your Amazon S3 bucket to Amazon Glacier on a weekly basis.

G. Use AWS Data Pipeline with multi-factor authentication to securely deploy code from the Amazon S3 bucket to your Amazon EC2 instances.

Answer: BD

Explanation:

You can add another layer of protection by enabling MFA Delete on a versioned bucket. Once you do so, you must provide your AWS account's access keys and a

valid code from the account's MFA device in order to permanently delete an object version or suspend or reactivate versioning on the bucket.

For more information on MFA please refer to the below link: <https://aws.amazon.com/blogs/security/securing-access-to-aws-using-mfa-part-3/>

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 IAM roles. For more information on Roles for EC2 please refer to the below link: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

Option A is invalid because this will not address either the integrity or security concern completely. Option C is invalid because user credentials should never be used in EC2 instances to access AWS resources.

Option D and F are invalid because AWS Pipeline is an unnecessary overhead when you already have inbuilt controls to manage security for S3.

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 CloudWatch Logs LogGroup and send the operating system and application logs of interest using the CloudWatch 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 are doing a load testing exercise on your application hosted on AWS. While testing your Amazon RDS MySQL DB instance, you notice that when you hit 100% CPU utilization on it, your application becomes non-responsive. Your application is read-heavy. What are methods to scale your data tier to meet the application's needs? Choose three answers from the options given below

- A. Add Amazon RDS DB read replicas, and have your application direct read queries to them.
- B. Add your Amazon RDS DB instance to an Auto Scaling group and configure your Cloud Watch metric based on CPU utilization.
- C. Use an Amazon SQS queue to throttle data going to the Amazon RDS DB instance.
- D. Use ElastiCache in front of your Amazon RDS DB to cache common queries.
- E. Shard your data set among multiple Amazon RDS DB instances.
- F. Enable Multi-AZ for your Amazon RDS DB instance.

Answer: ADE

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

For more information on Read Replica's please refer to the below link:

? <https://aws.amazon.com/rds/details/read-replicas/>

Sharding is a common concept to split data across multiple tables in a database For more information on sharding please refer to the below link:

<https://forums.aws.amazon.com/thread.jspa?messageID=203052>

Amazon ElastiCache is a web service that makes it easy to deploy, operate, and scale an in-memory data store or cache in the cloud. The service improves the performance of web applications by allowing you to retrieve information from fast, managed, in-memory data stores, instead of relying entirely on slower disk-based databases

Amazon ElastiCache is an in-memory key/value store that sits between your application and your database. Whenever your application requests data, it first makes the request to the ElastiCache cache. If the data exists in the cache and is current, ElastiCache returns the data to your application. If the data does not exist in the cache, or the data in the cache has expired, your application requests the data from your data store which returns the data to your application. Your application then writes the data received from the store to the cache so it can be more quickly retrieved next time it is requested. For more information on ElastiCache please refer to the below link:

<https://aws.amazon.com/elasticache/>

Option B is not an ideal way to scale a database

Option C is not ideal to store the data which would go into a database because of the message size Option F is invalid because Multi-AZ feature is only a failover option

NEW QUESTION 10

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

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

Answer: A

Explanation:

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

For more information Elastic Beanstalk managing worker environments, please visit the below URL:

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

NEW QUESTION 13

You are using CloudFormation to launch an EC2 instance and then configure an application after the instance is launched. You need the stack creation of the ELB and Auto Scaling to wait until the EC2 instance is launched and configured properly. How do you do this?

- A. It is not possible for the stack creation to wait until one service is created and launched
- B. Use the WaitCondition resource to hold the creation of the other dependent resources
- C. Use a CreationPolicy to wait for the creation of the other dependent resources >/
- D. Use the HoldCondition resource to hold the creation of the other dependent resources

Answer: C

Explanation:

When you provision an Amazon EC2 instance in an AWS CloudFormation stack, you might specify additional actions to configure the instance, such as install software packages or bootstrap applications. Normally, CloudFormation proceeds with stack creation after the instance has been successfully created. However, you can use a CreationPolicy so that CloudFormation proceeds with stack creation only after your configuration actions are done. That way you'll know your applications are ready to go after stack creation succeeds.

A CreationPolicy instructs CloudFormation to wait on an instance until CloudFormation receives the specified number of signals

Option A is invalid because this is possible

Option B is invalid because this is used make AWS CloudFormation pause the creation of a stack and wait for a signal before it continues to create the stack
 For more information on this, please visit the below URL:

- <https://aws.amazon.com/blogs/devops/use-a-creationpolicy-to-wait-for-on-instance-configurations/>

NEW QUESTION 18

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

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

Answer: A

Explanation:

Auto Scaling periodically performs health checks on the instances in your Auto Scaling group and identifies any instances that are unhealthy. You can configure Auto Scaling to determine the health status of an instance using Amazon EC2 status checks, Elastic Load Balancing health checks, or custom health checks. By default, Auto Scaling health checks use the results of the EC2 status checks to determine the health status of an instance. Auto Scaling marks an instance as unhealthy if its instance fails one or more of the status checks.

For more information monitoring in Autoscaling, please visit the below URL: <http://docs.aws.amazon.com/autoscaling/latest/userguide/monitoring-features.html>

NEW QUESTION 20

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

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

Answer: B

Explanation:

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

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

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

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

NEW QUESTION 24

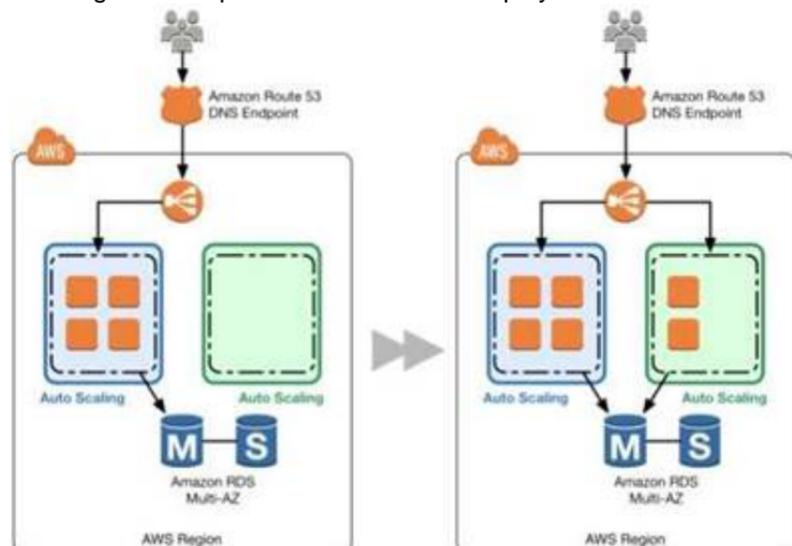
After a daily scrum with your development teams, you've agreed that using Blue/Green style deployments would benefit the team. Which technique should you use to deliver this new requirement?

- 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 Auto scaling 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 27

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 EC2 instances to get those instances which need to be removed from the configuration management system.

NEW QUESTION 30

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

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

Answer: A

Explanation:

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

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

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

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

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



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

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

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 Glacier
- F. Ensure that the operating system shutdown procedure triggers a logs transmission when the Amazon EC2 instance is stopped/terminated
- 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 bucket
- I. Ensure that the operating system shutdown procedure triggers a logs transmission when the Amazon EC2 instance is stopped/terminated
- 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 process
- 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 EC2 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/QuickStartEC2Instance.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 34

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

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

Answer: B

Explanation:

Custom resources enable you to write custom provisioning logic in templates that AWS 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.

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 37

You work for a startup that has developed a new photo-sharing application for mobile devices. Over recent months your application has increased in popularity; this has resulted in a decrease in the performance of the application due to the increased load. Your application has a two-tier architecture that is composed of an Auto Scaling PHP application tier and a MySQL RDS instance initially deployed with AWS CloudFormation. Your Auto Scaling group has a min value of 4 and a max value of 8. The desired capacity is now at 8 because of the high CPU utilization of the instances. After some analysis, you are confident that the performance issues stem from a constraint in CPU capacity, although memory utilization remains low. You therefore decide to move from the general-purpose M3 instances to the compute-optimized C3 instances. How would you deploy this change while minimizing any interruption to your end users?

- A. Sign into the AWS Management Console, copy the old launch configuration, and create a new launch configuration that specifies the C3 instance
- B. Update the Auto Scaling group with the new launch configuration
- C. Auto Scaling will then update the instance type of all running instances.
- D. Sign into the AWS Management Console, and update the existing launch configuration with the new C3 instance type
- E. Add an UpdatePolicy attribute to your Auto Scaling group that specifies AutoScalingRollingUpdate.
- F. Update the launch configuration specified in the AWS CloudFormation template with the new C3 instance type
- G. Run a stack update with the new template
- H. Auto Scaling will then update the instances with the new instance type.
- I. Update the launch configuration specified in the AWS CloudFormation template with the new C3 instance type
- J. Also add an UpdatePolicy attribute to your Auto Scaling group that specifies AutoScalingRollingUpdate
- K. Run a stack update with the new template.

Answer: D

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 rolling updates, please visit the below link:

- <https://aws.amazon.com/premiumsupport/knowledge-center/auto-scaling-group-rolling-updates/>

NEW QUESTION 41

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

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

Answer: AB

Explanation:

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

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

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

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

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

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

NEW QUESTION 44

You are using a configuration management system to manage your Amazon EC2 instances. On your Amazon EC2 Instances, you want to store credentials for connecting to an Amazon RDS MYSQL DB instance. How should you securely store these credentials?

- A. Give the Amazon EC2 instances an 1AM role that allows read access to a private Amazon S3 bucket
- B. Store a file with database credentials in the Amazon S3 bucket
- C. Have your configuration management system pull the file from the bucket when it is needed.
- D. Launch an Amazon EC2 instance and use the configuration management system to bootstrap the instance with the Amazon RDS DB credential
- E. Create an AMI from this instance.
- F. Store the Amazon RDS DB credentials in Amazon EC2 user data
- G. Import the credentials into the Instance on boot.
- H. Assign an 1AM role to your Amazon EC2 instance, and use this 1AM role to access the Amazon RDS DB from your Amazon EC2 instances.

Answer: D

Explanation:

Creating and Using an IAM Policy for IAM Database Access

To allow an IAM user or role to connect to your DB instance or DB cluster, you must create an IAM policy. After that you attach the policy to an IAM user or role.

Note
 To learn more about IAM policies, see Authentication and Access Control for Amazon RDS.

The following example policy allows an IAM user to connect to a DB instance using IAM database authentication.

```

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

```

Important

Don't confuse the rds-db: prefix with other Amazon RDS action prefixes that begin with rds:. You use the rds-db: prefix and the rds-db:connect action only for IAM database authentication. They aren't valid in any other context.

1AM Database Authentication for MySQL and Amazon Aurora

With Amazon RDS for MySQL or Aurora with MySQL compatibility, you can authenticate to your DB instance or DB cluster using AWS Identity and Access Management (IAM) database authentication. With this authentication method, you don't need to use a password when you connect to a DB instance. Instead, you use an authentication token.

An authentication token is a unique string of characters that Amazon RDS generates on request. Authentication tokens are generated using AWS Signature Version 4. Each token has a lifetime of 15 minutes. You don't need to store user credentials in the database, because authentication is managed externally using IAM. You can also still use standard database authentication.

IAM database authentication provides the following benefits:

- Network traffic to and from the database is encrypted using Secure Sockets Layer (SSL).
- You can use IAM to centrally manage access to your database resources, instead of managing access individually on each DB instance or DB cluster.
- For applications running on Amazon EC2, you can use EC2 instance profile credentials to access the database instead of a password, for greater security.

For more information please refer to the below document link from AWS

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.IAMDBAuth.html>

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.IAMDBAuth.IAMPolicy.html>

You can use roles to delegate access to users, applications, or services that don't normally have access to your AWS resources. For example, you might want to grant users in your AWS account access to resources they don't usually have, or grant users in one AWS account access to resources in another account. Or you might want to allow a mobile app to use AWS resources, but not want to embed AWS keys within the app (where they can be difficult to rotate and where users can potentially extract them). Sometimes you want to give AWS access to users who already have identities defined outside of AWS, such as in your corporate directory. Or, you might want to grant access to your account to third parties so that they can perform an audit on your resources. For more information on IAM Roles, please refer to the below document link: from AWS

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

NEW QUESTION 46

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

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

Answer: B

Explanation:

Ensure first that the cloudformation template is updated with the new instance type.

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

an update to the Cloud Formation stack occurs. A common approach to updating an Auto Scaling group is to perform a rolling update, which is done by specifying the AutoScalingRollingUpdate policy. This retains the same Auto Scaling group and replaces old instances with new ones, according to the parameters specified.

Option A is invalid because this will cause an interruption to the users.

Option C is partially correct, but it does not have all the steps as mentioned in option B.

Option D is partially correct, but we need the AutoScalingRollingUpdate attribute to ensure a rolling update is performed.

For more information on AutoScaling Rolling updates please refer to the below link:

- <https://aws.amazon.com/premiumsupport/knowledge-center/auto-scaling-group-rolling-updates/>

NEW QUESTION 48

What is web identity federation?

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

Answer: B

Explanation:

With web identity federation, you don't need to create custom sign-in code or manage your own user identities. Instead, users of your app can sign in using a well-known identity provider (IdP) — such as Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible IdP, receive an authentication token, and then exchange that token for temporary security credentials in AWS that map to an IAM role with permissions to use the resources in your AWS account. Using an IdP helps you keep your AWS account secure, because you don't have to embed and distribute long-term security credentials with your application. For more information on Web Identity federation please refer to the below link:

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

NEW QUESTION 53

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 f, inclusive.

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 have 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 56

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

- A. Copy all log files into AWS S3 using a cron job on each instance
- B. Use an S3 Notification Configuration on the PutBucket event and publish events to AWS Lambda
- C. Use the Lambda to analyze logs as soon as they come in and flag issues.
- D. Begin using CloudWatch Logs on every service
- E. Stream all Log Groups into S3 object

- F. Use AWS EMR clusterjobs to perform adhoc MapReduce analysis and write new queries when needed.
- G. Copy all log files into AWS S3 using a cron job on each instanc
- H. Use an S3 Notification Configuration on the PutBucket event and publish events to AWS Kinesi
- I. Use Apache Spark on AWS EMR to perform at-scale stream processing queries on the log chunks and flag issues.
- J. Begin using CloudWatch Logs on every servic
- K. Stream all Log Groups into an AWS Elastic search Service Domain running Kibana 4 and perform log analysis on a search cluster.

Answer: D

Explanation:

Amazon 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 Cloud Watch so that you can go from raw data to actionable insights quickly. For more information on Elastic Search, please refer to the below link:
 • <https://aws.amazon.com/elasticsearch-service/>

NEW QUESTION 59

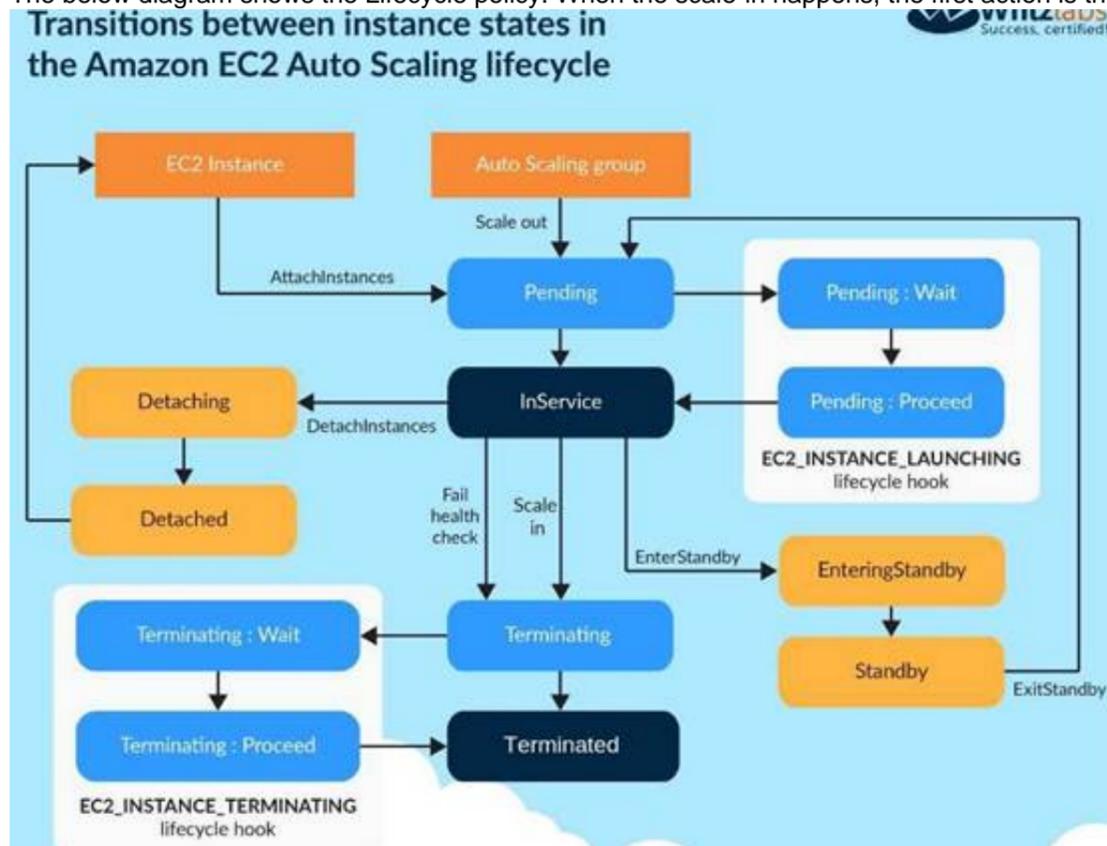
For AWS Auto Scaling, what is the first transition state an instance enters after leaving steady state when scaling in due to health check failure or decreased load?

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

Answer: A

Explanation:

The below diagram shows the Lifecycle policy. When the scale-in happens, the first action is the Terminating action.



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

NEW QUESTION 63

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

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

Answer: ACD

Explanation:

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

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

NEW QUESTION 67

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 Elastic Compute Cloud (Amazon EC2) 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 71

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 EC2 please refer to the below link: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

NEW QUESTION 76

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 77

You want to pass queue messages that are 1GB each. How should you achieve this?

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

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 78

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. AWS DynamoDB

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 83

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

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

Answer: A

Explanation:

Provisioned IOPS SSD should be used for critical business applications that require sustained IOPS performance, or more than 10,000 IOPS or 160 MiB/s of throughput per volume

This is ideal for Large database workloads, such as:

- MongoDB
- Cassandra
- MicrosoftSQL Server
- MySQL
- PostgreSQL
- Oracle

For more information on the various CBS Volume Types, please refer to the below link:

- <http://docs.aws.amazon.com/AWSSCC2/latest/UserGuide/CBSVolumeTypes.html>

NEW QUESTION 86

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 90

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 94

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 98

You are planning on using encrypted snapshots in the design of your AWS Infrastructure. Which of the following statements are true with regards to EBS Encryption

- A. Snapshotting an encrypted volume makes an encrypted snapshot; restoring an encrypted snapshot creates an encrypted volume when specified / requested.
- B. Snapshotting an encrypted volume makes an encrypted snapshot when specified / requested; restoring an encrypted snapshot creates an encrypted volume when specified / requested.
- C. Snapshotting an encrypted volume makes an encrypted snapshot; restoring an encrypted snapshot always creates an encrypted volume.
- D. Snapshotting an encrypted volume makes an encrypted snapshot when specified / requested; restoring an encrypted snapshot always creates an encrypted volume.

Answer: C

Explanation:

Amazon CBS encryption offers you a simple encryption solution for your CBS volumes without the need for you to build, maintain, and secure your own key management infrastructure. When you create an encrypted CBS volume and attach it to a supported instance type, the following types of data are encrypted:

- Data at rest inside the volume
- All data moving between the volume and the instance
- All snapshots created from the volume

Snapshots that are taken from encrypted volumes are automatically encrypted. Volumes that are created from encrypted snapshots are also automatically encrypted.

For more information on CBS encryption, please visit the below URL:

- <http://docs.aws.amazon.com/AWSC2/latest/UserGuide/CBSEncryption.html>

NEW QUESTION 102

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

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

Answer: B

Explanation:

The CBS Snapshots are stored in S3, so if you have an scripts which deploy EC2 Instances, the CBS volumes need to be constructed from snapshots stored in S3.

You can back up the data on your Amazon CBS volumes to Amazon S3 by taking point-in-time snapshots. Snapshots are incremental backups, which means that only the blocks on the device that have changed after your most recent snapshot are saved. This minimizes the time required to create the snapshot and saves on storage costs by not duplicating data. When you delete a snapshot, only the data unique to that snapshot is removed. Each snapshot contains all of the information needed to restore your data (from the moment when the snapshot was taken) to a new CBS volume. For more information on CBS Snapshots, please visit the below URL:

- <http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/CBSSnapshots.htm>

NEW QUESTION 104

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

- A. AWS CloudFormation
- B. AWS OpsWorks
- C. AWS ELB+ EC2 with CLI Push
- D. AWS Elastic Beanstalk

Answer: D

Explanation:

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

- <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/Welcome.html>

NEW QUESTION 105

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 110

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 115

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

- A. Use CloudTrail to log all the API calls, and then traverse the log files to locate the issue
- B. Use CloudWatch, monitor the CPU utilization to see the times when the CPU peaked

- C. Review the Elastic Load Balancer logs
- D. Create some custom Cloudwatch metrics which are pertinent to the key features of your application

Answer: D

Explanation:

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

NEW QUESTION 118

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 120

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

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

Answer: B

Explanation:

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

NEW QUESTION 125

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

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

Answer: ABD

Explanation:

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

- 1) Amazon Kinesis - Here data can be fed for real time analysis
- 2) Amazon S3 - You can use CloudWatch Logs to store your log data in highly durable storage such as S3.
- 3) Amazon Lambda - Lambda functions can be designed to work with Cloudwatch log. For more information on Cloudwatch Logs, please refer to the below link:

link:<http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html>

NEW QUESTION 126

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

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

Answer: C

Explanation:

With the Swap url feature, you can keep a version of your environment ready. And when you are ready to cut over, you can just use the swap url feature to switch over

to your new environment

For more information on swap url feature, please refer to the below link:

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

NEW QUESTION 128

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

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

Answer: AB

Explanation:

The Auto Scaling cooldown period is a configurable setting for your Auto Scaling group that helps to ensure that Auto Scaling doesn't launch or terminate additional instances before the previous scaling activity takes effect.

For more information on the cool down period, please refer to the below link:

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

Also it is better to monitor the application based on a key feature and then trigger the scale-in and scale-out feature accordingly. In the question, there is no mention of CPU or memory causing the issue.

NEW QUESTION 130

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 132

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

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 135

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 137

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 138

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 142

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 145

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 150

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 154

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 155

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

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

Answer: D

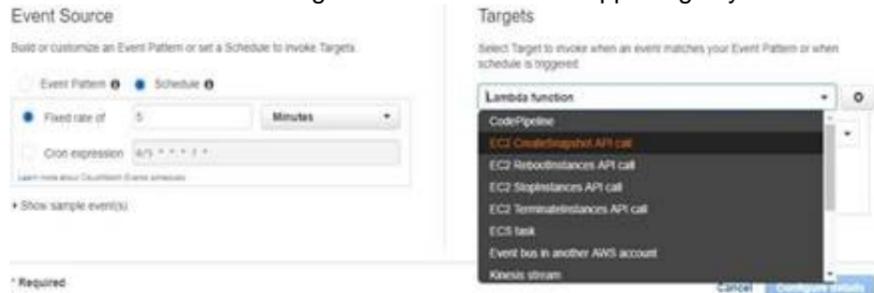
Explanation:

The best is to use the inbuilt service from Cloudwatch, as Cloud watch Events to automate the creation of CBS Snapshots. With Option A, you would be restricted to

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

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

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



The AWS Documentation mentions

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

that you can quickly set up, you can match events and route them to one or more target functions or streams. Cloud Watch Events becomes aware of operational

changes as they occur. Cloud Watch 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 159

You currently have an Autoscalinggroup that has the following settings Min capacity-2

Desired capacity - 2 Maximum capacity - 2

Your launch configuration has AMI'S which are based on the t2.micro instance type. The application running on these instances are now experiencing issues and you have identified that the solution is to change the instance type of the instances running in the Autoscaling Group.

Which of the below solutions will meet this demand.

- A. Change the Instance type in the current launch configuratio
- B. Change the Desired value of the Autoscaling Group to 4. Ensure the new instances are launched.
- C. Delete the current Launch configuratio
- D. Create a new launch configuration with the new instance type and add it to the Autoscaling Grou
- E. This will then launch the new instances.
- F. Make a copy the Launch configuratio
- G. Change the instance type in the new launch configuratio
- H. Attach that to the Autoscaling Group.Change the maximum and Desired size of the Autoscaling Group to 4. Once the new instances are launched, change the Desired and maximum size back to 2.
- I. Change the desired and maximum size of the Autoscaling Group to 4. Make a copy the Launch configuratio
- J. Change the instance type in the new launch configuratio
- K. Attach that to the Autoscaling Grou
- L. Change the maximum and Desired size of the Autoscaling Group to 2

Answer: C

Explanation:

You should make a copy of the launch configuration, add the new instance type. The change the Autoscaling Group to include the new instance type. Then change the Desired number of the Autoscaling Group to 4 so that instances of new instance type can be launched. Once launched, change the desired size back to 2, so that Autoscaling will delete the instances with the older configuration. Note that the assumption here is that the current instances are equally distributed across multiple AZ's because Autoscaling will first use the AZRebalance process to terminate instances.

Option A is invalid because you cannot make changes to an existing Launch configuration.

Option B is invalid because if you delete the existing launch configuration, then your application will not be available. You need to ensure a smooth deployment process.

Option D is invalid because you should change the desired size to 4 after attaching the new launch configuration.

For more information on Autoscaling Suspend and Resume, please visit the below URL: <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resume-processes.html>

NEW QUESTION 160

You are currently using Elastic Beanstalk to host your production environment. You need to rollout updates to your application hosted on this environment. This is a critical application which is why there is a requirement that the rollback, if required, should be carried out with the least amount of downtime. Which of the following deployment strategies would ideally help achieve this purpose

- A. Create a Cloudformation template with the same resources as those in the Elastic beanstalk environmen
- B. If the deployment fails, deploy the Cloudformation template.
- C. Use Rolling updates in Elastic Beanstalk so that if the deployment fails, the rolling updates feature would roll back to the last deployment.
- D. Create another parallel environment in elastic beanstal
- E. Use the Swap URL feature.
- F. Create another parallel environment in elastic beanstal
- G. Create a new Route53 Domain name for the new environment and release that url to the users.

Answer: C

Explanation:

Since the requirement is to have the least amount of downtime, the ideal way is to create a blue green deployment environment and then use the Swap URL feature

to swap environments for the new deployment and then do the swap back, incase the deployment fails.

The AWS Documentation mentions the following on the SWAP url feature of Elastic Beanstalk

Because Elastic Beanstalk performs an in-place update when you update your application versions, your application may become unavailable to users for a short period of time. It is possible to avoid this downtime by performing a blue/green deployment, where you deploy the new version to a separate environment, and then swap CNAMCs of the two environments to redirect traffic to the new version instantly.

NEW QUESTION 163

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

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

Answer: C

Explanation:

The AWS Documentation mentions the following

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

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

NEW QUESTION 166

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

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

Answer: C

Explanation:

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

For more information on IAM please see the below link:

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

NEW QUESTION 168

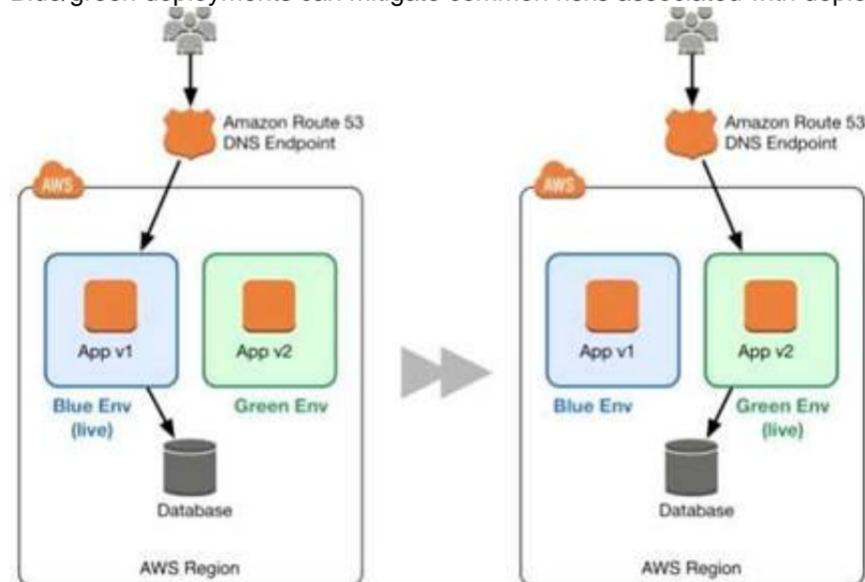
A company has developed a Ruby on Rails content management platform. Currently, OpsWorks with several stacks for dev, staging, and production is being used to deploy and manage the application. Now the company wants to start using Python instead of Ruby. How should the company manage the new deployment?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

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

Answer: B

Explanation:

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



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

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

NEW QUESTION 169

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

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

Answer: A

Explanation:

AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account. With CloudTrail, you can log, continuously monitor, and retain events related to API calls across your AWS infrastructure. CloudTrail provides a history of AWS API calls for your account, including API calls made through the AWS Management Console, AWS SDKs, command line tools, and other AWS services. This history simplifies security analysis, resource change tracking, and troubleshooting.

You can design cloudtrail to send all logs to a central S3 bucket. For more information on cloudtrail, please visit the below URL:

? <https://aws.amazon.com/cloudtrail/>

NEW QUESTION 170

Your company is getting ready to do a major public announcement of a social media site on AWS. The website is running on EC2 instances deployed across multiple Availability Zones with a Multi-AZ RDS MySQL Extra Large DB Instance. The site performs a high number of small reads and writes per second and relies on an eventual consistency model. After comprehensive tests you discover that there is read contention on RDS MySQL. Which are the best approaches to meet these requirements? Choose 2 answers from the options below

- A. Deploy ElasticCache in-memory cache running in each availability zone
- B. Implement sharding to distribute load to multiple RDS MySQL instances
- C. Increase the RDS MySQL Instance size and Implement provisioned IOPS
- D. Add an RDS MySQL read replica in each availability zone

Answer: AD

Explanation:

Implement Read Replicas and Elastic Cache

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

For more information on Read Replica's, please visit the below link:

- <https://aws.amazon.com/rds/details/read-replicas/>

Amazon ElastiCache is a web service that makes it easy to deploy, operate, and scale an in-memory data store or cache in the cloud. The service improves the performance of web applications by allowing you to retrieve information from fast, managed, in-memory data stores, instead of relying entirely on slower disk-based databases.

For more information on Amazon ElastiCache, please visit the below link:

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

NEW QUESTION 172

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

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

Answer: B

Explanation:

The AWS documentation mentions

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

When your usage exceeds these amounts, AWS sends you an

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

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

NEW QUESTION 174

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

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

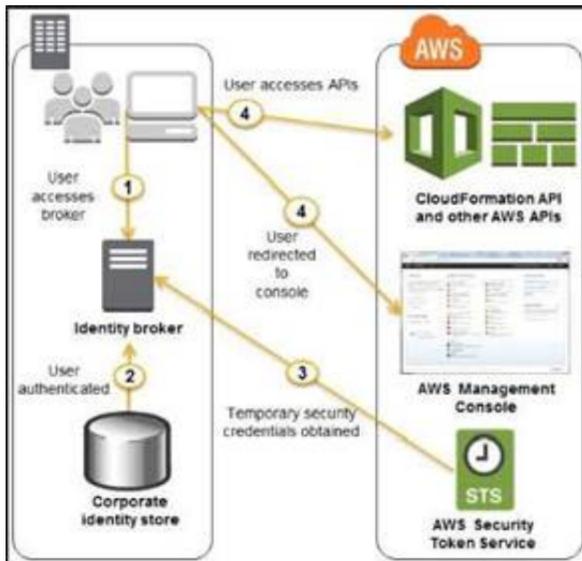
Answer: BC

Explanation:

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

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

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



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

NEW QUESTION 175

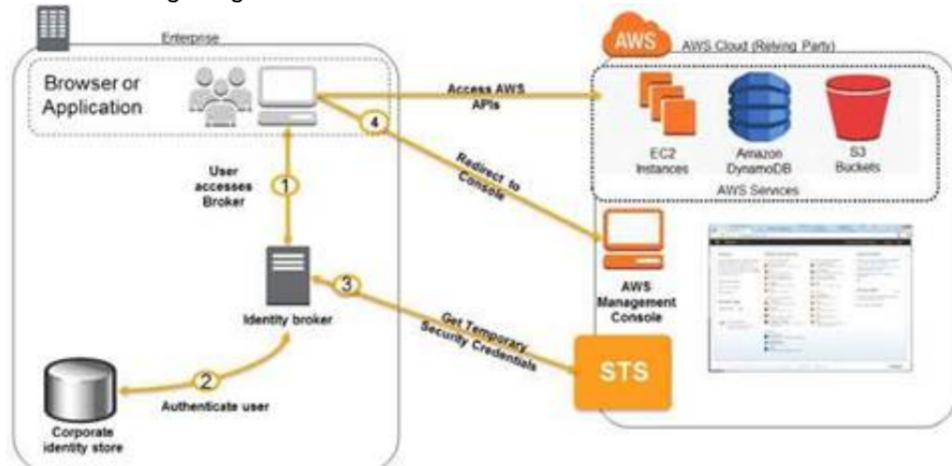
Which of the following will you need to consider so you can set up a solution that incorporates single sign-on from your corporate AD or LDAP directory and restricts access for each user to a designated user folder in a bucket? Choose 3 Answers from the options below

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

Answer: ABD

Explanation:

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



For more information on federated access, please visit the below link: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_common-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 178

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 182

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

- A. Pipeline definition
- B. Task Runner
- C. Task History

D. Workflow Runner

Answer: AB

Explanation:

The AWS Documentation mentions the following on AWS Pipeline

The following components of AWS Data Pipeline work together to manage your data: A pipeline definition specifies the business logic of your data management. A pipeline schedules and runs tasks. You upload your pipeline definition to the pipeline, and then activate the pipeline. You can edit the pipeline definition for a running pipeline and activate the pipeline again for it to take effect. You can deactivate the pipeline, modify a data source, and then activate the pipeline again. When you are finished with your pipeline, you can delete it.

Task Runner polls for tasks and then performs those tasks. For example, Task Runner could copy log files to Amazon S3 and launch Amazon EMR clusters. Task Runner is installed and runs automatically on resources created by your pipeline definitions. You can write a custom task runner application, or you can use the Task Runner application that is provided by AWS Data Pipeline.

For more information on AWS Pipeline, please visit the link: <http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/what-is-datapipeline.html>

NEW QUESTION 185

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

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

Answer: B

Explanation:

The best way to get notified when the images get sent to the bucket is the S3 Events. Here you don't need to provision infrastructure beforehand, and also since the source of S3 provides Event Handling, this should be used.

The AWS Documentation mentions the following on AWS S3 and Lambda

Amazon S3 can publish events (for example, when an object is created in a bucket) to AWS Lambda and invoke your Lambda function by passing the event data as a parameter. This integration enables

you to write Lambda functions that process Amazon S3 events. In Amazon S3, you add bucket notification configuration that identifies the type of event that you want Amazon S3 to publish and the Lambda function that you want to invoke.

For more information on AWS Lambda and S3 Events, please visit the link:

- <http://docs.aws.amazon.com/lambda/latest/dg/with-s3.html>

NEW QUESTION 187

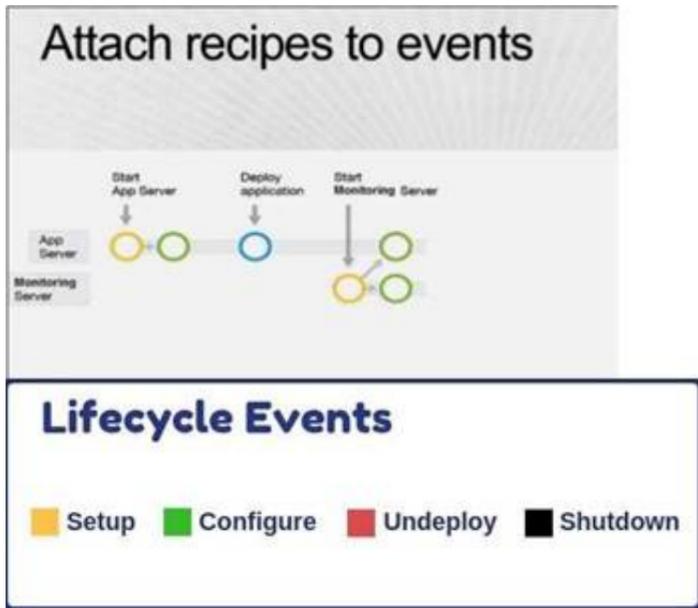
Which of the following are Lifecycle events available in Opswork? Choose 3 answers from the options below

- 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 190

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

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

Answer: B

Explanation:

The AWS documentation mentions the following
 You can manage Amazon SQS messages with Amazon S3. This is especially useful for storing and consuming messages with a message size of up to 2 GB. To manage Amazon SQS messages with Amazon S3, use the Amazon SQS Extended Client Library for Java. Specifically, you use this library to:
 Specify whether messages are always stored in Amazon S3 or only when a message's size exceeds 256 KB.
 Send a message that references a single message object stored in an Amazon S3 bucket. Get the corresponding message object from an Amazon S3 bucket.
 Delete the corresponding message object from an Amazon S3 bucket. For more information on SQS and sending larger messages please visit the link

NEW QUESTION 191

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

- A. Createseparate dashboards in every region
- B. Register!nstances running on different regions to CloudWatch
- C. Haveone single dashboard to report metrics to CloudWatch from different region
- D. Thisis not possible

Answer: C

Explanation:

You can monitor AWS resources in multiple regions using a single Cloud Watch dashboard. For example, you can create a dashboard that shows CPU utilization for an EC2 instance located in the us-west-2 region with your billing metrics, which are located in the us-east-1 region.
 For more information on Cloudwatch dashboard, please refer to the below url
http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/cross_region_dashboard.html

NEW QUESTION 192

Your company is getting ready to do a major public announcement of a social media site on AWS. The website is running on EC2 instances deployed across multiple Availability Zones with a Multi-AZ RDS MySQL Extra Large DB Instance. The site performs a high number of small reads and writes per second and relies on an eventual consistency model. After comprehensive tests you discover that there is read contention on RDS MySQL. Which are the best approaches to meet these requirements? Choose 2 answers from the options below

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

Answer: AD

Explanation:

Implement Read Replicas and Elastic Cache
 Amazon RDS Read Replicas provide enhanced performance and durability for database (DB) instances. This replication feature makes it easy to elastically scale out beyond the capacity constraints of a single DB Instance for read-heavy database workloads. You can create one or more replicas of a given source DB

Instance and serve high-volume application read traffic from multiple copies of your data, thereby increasing aggregate read throughput.

For more information on Read Replica's, please visit the below link

- <https://aws.amazon.com/rds/details/read-replicas/>

Amazon OastlCache is a web service that makes it easy to deploy, operate, and scale an in-memory data store or cache in the cloud. The service improves the performance of web applications by allowing you to retrieve information from fast, managed, in-memory data stores, instead of relying entirely on slower disk-based databases.

For more information on Amazon OastlCache, please visit the below link

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

NEW QUESTION 194

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

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

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 198

Your company owns multiple AWS accounts. There is currently one development and one production account. You need to grant access to the development team to an S3 bucket in the production account. How can you achieve this?

- A. Create an IAM user in the Production account that allows users from the Development account (the trusted account) to access the S3 bucket in the Production account.
- B. When creating the role, define the Development account as a trusted entity and specify a permissions policy that allows trusted users to update the S3 bucket.
- C. Use web identity federation with a third-party identity provider with AWS STS to grant temporary credentials and membership into the production IAM user.
- D. Create an IAM cross account role in the Production account that allows users from the Development account to access the S3 bucket in the Production account.

Answer: D

Explanation:

The AWS Documentation mentions the following on cross account roles

You can use AWS Identity and Access Management (IAM) roles and AWS Security Token Service (STS) to set up cross-account access between AWS accounts. When you assume an IAM role in another AWS account to obtain cross-account access to services and resources in that account, AWS CloudTrail logs the cross-account activity. For more information on Cross account roles, please visit the below URL

- http://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 203

You work at a company that makes use of AWS resources. One of the key security policies is to ensure that all data is encrypted both at rest and in transit. Which of the following is not a right implementation which aligns to this policy?

- A. Using S3 Server Side Encryption (SSE) to store the information
- B. Enable SSL termination on the ELB
- C. Enabling Proxy Protocol ID- Enabling sticky sessions on your load balancer

Answer: B

Explanation:

Please note the keyword "NOT" in the question.

Option A is incorrect. Enabling S3 SSE encryption helps the encryption of data at rest in S3. So Option A is invalid.

Option B is correct. If you disable SSL termination on the ELB the traffic will be encrypted all the way to the backend. SSL termination allows encrypted traffic between the client

and the ELB but cause traffic to be unencrypted between the ELB and the backend (presumably EC2 or ECS/Task, etc.)

If SSL is not terminated on the ELB you must use Layer A to have traffic encrypted all the way.

Sticky sessions are not supported with Layer A (TCP endpoint). Thus option D "Enabling sticky sessions on your load balancer" can't be used and is the right answer

For more information on sticky sessions, please visit the below URL <https://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-sticky-sessions.html>

Requirements

- An HTTP/HTTPS load balancer.
- At least one healthy instance in each Availability Zone.
- At least one healthy instance in each Availability Zone.

If you don't want the load balancer to handle the SSL termination (known as SSL offloading), you can use TCP for both the front-end and back-end connections, and deploy certificates on the registered instances handling requests.

For more information on elb-listener-config, please visit the below

- <https://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-listener-config.html> If the front-end connection uses TCP or SSL, then your back-end connections can use either TCP or SSL. Note: You can use an HTTPS listener and still use SSL on the backend but the ELB must terminate, decrypt and re-encrypt. This is slower and less secure than using the same encryption all the way to the backend.. It also breaks the question requirement of having all data encrypted in transit since it forces the ELB to decrypt. Proxy protocol is used to provide a secure transport connection hence Option C is also incorrect. For more information on SSL Listeners for your load balancer, please visit the below URL
<http://docsaws.amazon.com/elasticloadbalancing/latest/classic/elb-https-load-balancers.html>
<https://aws.amazon.com/blogs/aws/elastic-load-balancer-support-for-ssl-termination/>

NEW QUESTION 204

.....

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