

DOP-C01 Dumps

AWS Certified DevOps Engineer- Professional

<https://www.certleader.com/DOP-C01-dumps.html>



NEW QUESTION 1

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 2

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

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

Answer: B

Explanation:

Elastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Each log contains information such as the

time the request was received, the client's IP address, latencies, request paths, and server responses.

You can use these access logs to analyze traffic patterns and to troubleshoot issues.

Option A is invalid because this service will monitor all AWS services. Option C and D are invalid since CLB already provides a logging feature.

For more information on ELB access logs, please refer to the below document link: from AWS

<http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/access-log-collection.html>

NEW QUESTION 3

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 4

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

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

NEW QUESTION 5

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

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

Answer: C

Explanation:

You can use Cloud Watch Logs to monitor applications and systems using log data. For example, CloudWatch Logs can track the number of errors that occur in your application logs and send you a notification whenever the rate of errors exceeds a threshold you specify. CloudWatch Logs uses your log data for monitoring; so, no code changes are required.

Option A and B are invalid because Autoscaling policies are not designed for these purposes. Option D is invalid because you use Cloudwatch Logs Agent and not the monitoring agent. For more information on Cloudwatch logs, please refer to the below link:

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

NEW QUESTION 6

You have the following application to be setup in AWS

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

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

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

Answer: D

Explanation:

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

the applications use. Instead of creating and distributing your AWS credentials For more information on IAM Roles please refer to the below link:

<http://docs.aws.amazon.com/AWSC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

NEW QUESTION 7

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

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

Answer: A

Explanation:

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

other templates. For more information on CloudFormation best practises please refer to the below link:

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

NEW QUESTION 8

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 the database that it accesses. 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 9

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 10

You have a multi-docker environment that you want to deploy to AWS. Which of the following configuration files can be used to deploy a set of Docker containers as an Elastic Beanstalk application?

- A. `Dockerfile`
- B. `.ebextensions`
- C. `Dockerfile`
- D. `Dockerfile`

Answer: A

Explanation:

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

NEW QUESTION 10

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

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

Answer: B

Explanation:

During the AMI-creation process, Amazon EC2 creates snapshots of your instance's root volume and any other EBS volumes attached to your instance. New EBS volumes receive their maximum performance the moment that they are available and do not require initialization (formerly known as pre-warming). However, storage blocks on volumes that were restored from snapshots must be initialized (pulled down from Amazon S3 and written to the volume) before you can access the block. This preliminary action takes time and can cause a significant increase in the latency of an I/O operation the first time each block is accessed. For most applications, amortizing this cost over the lifetime of the volume is acceptable. Option A is invalid because block sizes are predetermined and should not be randomly selected. Option C is invalid because this is part of continuous integration and hence volumes can be destroyed after the test and hence there should not be snapshots created unnecessarily. Option D is invalid because the encryption is a security feature and not part of load tests normally. For more information on EBS initialization please refer to the below link:

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

NEW QUESTION 15

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/as-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. It's 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 21

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

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

Answer: A

Explanation:

Elastic Beanstalk supports the deployment of web applications from Docker containers. With Docker containers, you can define your own runtime environment. You can choose your own platform, programming language, and any application dependencies (such as package managers or tools), that aren't supported by other platforms. Docker containers are self-contained and include all the configuration information and software your web application requires to run. By using Docker with Elastic Beanstalk, you have an infrastructure that automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring.

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

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

NEW QUESTION 25

You are using Elastic Beanstalk to manage your application. You have a SQL script that needs to only be executed once per deployment no matter how many EC2 instances you have running. How can you do this?

- A. Use a "Container command" within an Elastic Beanstalk configuration file to execute the script, ensuring that the "leader only" flag is set to false.
- B. Use Elastic Beanstalk version and a configuration file to execute the script, ensuring that the "leader only" flag is set to true.
- C. Use a "Container command" within an Elastic Beanstalk configuration file to execute the script, ensuring that the "leader only" flag is set to true.
- D. Use a "leader command" within an Elastic Beanstalk configuration file to execute the script, ensuring that the "container only" flag is set to true.

Answer: C

Explanation:

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

You can use `leader_only` to only run the command on a single instance, or configure a test to only run the command when a test command evaluates to true.

Leader-only container commands are only executed during environment creation and deployments, while other commands and server customization operations are performed every time an instance is provisioned or updated. Leader-only container commands are not executed due to launch configuration changes, such as a change in the AMI Id or instance type. For more information on customizing containers, please visit the below URL:

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/customize-containers-ec2.html>

NEW QUESTION 30

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 CloudFormation stack occurs. A common approach to updating an Auto Scaling group is to perform a rolling update, which is done by specifying the `AutoScalingRollingUpdate` policy. This retains the same Auto Scaling group and replaces old instances with new ones, according to the parameters specified. 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 33

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

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 34

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

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

Answer: C

Explanation:

You can save your environment's configuration as an object in Amazon S3 that can be applied to other environments during environment creation, or applied to a running environment. Saved configurations are YAML formatted templates that define an environment's platform configuration, tier, configuration option settings, and tags.

For more information on Saved Configurations please refer to the below link:

- <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environment-configuration-savedconfig.html>

NEW QUESTION 38

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 41

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 IAM 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 IAM role to your Amazon EC2 instance, and use this IAM 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:rdbuser: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 45

You have been asked to de-risk deployments at your company. Specifically, the CEO is concerned about outages that occur because of accidental inconsistencies between Staging and Production, which sometimes cause unexpected behaviors in Production even when Staging tests pass. You already use Docker to get high consistency between Staging and Production for the application environment on your EC2 instances. How do you further de-risk the rest of the execution environment, since in AWS, there are many service components you may use beyond EC2 virtual machines?

- A. Develop models of your entire cloud system in CloudFormation
- B. Use this model in Staging and Production to achieve greater parity
- C. */
- D. Use AWS Config to force the Staging and Production stacks to have configuration parity
- E. Any differences will be detected for you so you are aware of risks.
- F. Use AMIs to ensure the whole machine, including the kernel of the virtual machines, is consistent, since Docker uses Linux Container (LXC) technology, and we need to make sure the container environment is consistent.
- G. Use AWS ECS and Docker clusterin
- H. This will make sure that the AMIs and machine sizes are the same across both environments.

Answer: A

Explanation:

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

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

NEW QUESTION 46

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

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

Answer: A

Explanation:

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

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

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

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

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

NEW QUESTION 50

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 51


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/AutoScalingGroupLifecycle.html>

NEW QUESTION 53

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 55

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 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>

NEW QUESTION 56

Which of these is not an intrinsic function in AWS CloudFormation?

- A. Fn::Equals
- B. Fn::If
- C. Fn::Not
- D. Fn::Parse

Answer: D

Explanation:

You can use intrinsic functions, such as Fn::If, Fn::Cqals, and Fn::Not, to conditionally create stack resources. These conditions are evaluated based on input parameters that you declare when you create or update a stack. After you define all your conditions, you can associate them with resources or resource properties in the Resources and Outputs sections of a template.

For more information on Cloud Formation template functions, please refer to the URL:

- <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference.html> and
- <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-conditions.html>

NEW QUESTION 59

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 61

You are designing a service that aggregates clickstream data in batch and delivers reports to subscribers via email only once per week. Data is extremely spikey, geographically distributed, high- scale, and unpredictable. How should you design this system?

- A. Use a large RedShift cluster to perform the analysis, and a fleet of Lambdas to perform recordinserts into the RedShift table
- B. Lambda will scale rapidly enough for the traffic spikes.
- C. Use a CloudFront distribution with access log delivery to S3. Clicks should be recorded as querystring GETs to the distributio
- D. Reports are built and sent by periodically running EMRjobs over the access logs in S3.C Use API Gateway invoking Lambdas which PutRecords into Kinesis, and EMR running Spark performing GetRecords on Kinesis to scale with spike
- E. Spark on EMR outputs the analysis to S3, which are sent out via email.D- Use AWS Elasticsearch service and EC2 Auto Scaling group
- F. The Autoscaling groups scale based on click throughput and stream into the Elasticsearch domain, which is also scalabl
- G. Use Kibana to generate reports periodically.

Answer: B

Explanation:

When you look at building reports or analyzing data from a large data set, you need to consider CMR because this service is built on the Hadoop framework which is used to processes large data sets.

The ideal approach to getting data onto CMR is to use S3. Since the Data is extremely spikey and geographically distributed, using edge locations via Cloudfront distributions is the best way to fetch the data.

Option A is invalid because RedShift is more of a petabyte storage cluster.

Option C is invalid because having both Kinesis and CMR for the job analysis is redundant. Option D is invalid because Elastic Search is not an option for

processing records.

For more information on Amazon CMR, please visit the below URL:

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

NEW QUESTION 64

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

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

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 67

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

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


Answer: C

Explanation:

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

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

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



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

NEW QUESTION 70

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/AWSCC2/latest/UserGuide/CBSVolumeTypes.html>

NEW QUESTION 72

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

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

Which services would suffice the above requirements

- A. AWS Config, CloudTrail, 1AM Credential Reports
- B. CloudTrail, 1AM Credential Reports, AWS Config
- C. CloudTrail, AWS Config, 1AM Credential Reports
- D. AWS Config, 1AM Credential Reports, CloudTrail

Answer: C

Explanation:

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

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

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

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

You can generate and download a credential report that lists all users in your account and the status of their various credentials, including passwords, access keys, and MFA devices. You can get a credential report from the AWS Management Console, the AWS SDKs and Command Line Tools, or the 1AM API. For more information on Credentials Report, please visit the below URL:

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

NEW QUESTION 74

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

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

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 76

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 78

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

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

Answer: A

Explanation:

Jumbo frames allow more than 1500 bytes of data by increasing the payload size per packet, and thus increasing the percentage of the packet that is not packet overhead. Fewer packets are needed to send the same amount of usable data. However, outside of a given AWS region (CC2-Classical), a single VPC, or a VPC peering

connection, you will experience a maximum path of 1500 MTU. VPN connections and traffic sent over an Internet gateway are limited to 1500 MTU. If packets are over

1500 bytes, they are fragmented, or they are dropped if the Don't Fragment flag is set in the IP header.

For more information on Jumbo Frames, please visit the below URL:

http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/network_mtu.htm#jumbo_frame_instances

NEW QUESTION 81

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

- A. Use OpsWorks Stacks with three layers to model the layering in your stack.
- B. Use CloudFormation Nested Stack Templates, with three child stacks to represent the three logical layers of your cloud.
- C. Use AWS Config to declare a configuration set that AWS should roll out to your cloud.
- D. Use Elastic Beanstalk Linked Applications, passing the important DNS entries between layers using the metadata interface.

Answer: B

Explanation:

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

For more information on nested stacks, please visit the below URL:

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

The query is, how you can automate a stack over the period of time, when changes are required, without recreating the stack.

The function of Nested Stacks are to reuse Common Template Patterns.

For example, assume that you have a load balancer configuration that you use for most of your stacks. Instead of copying and pasting the same configurations into your templates, you can create a dedicated template for the load balancer. Then, you just use the resource to reference that template from within other templates.

Yet another example is if you have a launch configuration with certain specific configuration and you need to change the instance size only in the production environment and to leave it as it is in the development environment.

AWS also recommends that updates to nested stacks are run from the parent stack.

When you apply template changes to update a top-level stack, AWS CloudFormation updates the top-level stack and initiates an update to its nested stacks. AWS CloudFormation updates the resources of modified nested stacks, but does not update the resources of unmodified nested stacks.

NEW QUESTION 84

Which of the following tools does not directly support AWS OpsWorks, for monitoring your stacks?

- A. AWSConfig
- B. Amazon CloudWatch Metrics
- C. AWS CloudTrail
- D. Amazon CloudWatch Logs

Answer: A

Explanation:

You can monitor your stacks in the following ways.

- AWS OpsWorks Stacks uses Amazon CloudWatch to provide thirteen custom metrics with detailed monitoring for each instance in the stack.
- AWS OpsWorks Stacks integrates with AWS CloudTrail to log every AWS OpsWorks Stacks API call and store the data in an Amazon S3 bucket.
- You can use Amazon CloudWatch Logs to monitor your stack's system, application, and custom logs.

For more information on Opswork monitoring, please visit the below URL:

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

NEW QUESTION 88

Your company wants to understand where cost is coming from in the company's production AWS account. There are a number of applications and services running at any given time. Without expending too much initial development time, how best can you give the business a good understanding of which applications

cost the most per month to operate?

- A. Create an automation script which periodically creates AWS Support tickets requesting detailed intra-month information about your bill.
- B. Use custom CloudWatch Metrics in your system, and put a metric data point whenever cost is incurred.
- C. Use AWS Cost Allocation Tagging for all resources which support it.
- D. Use the Cost Explorer to analyze costs throughout the month.
- E. Use the AWS Price API and constantly running resource inventory scripts to calculate total price based on multiplication of consumed resources over time.

Answer: C

Explanation:

A tag is a label that you or AWS assigns to an AWS resource. Each tag consists of a key and a value. A key can have more than one value. You can use tags to organize your resources, and cost allocation tags to track your AWS costs on a detailed level. After you activate cost allocation tags, AWS uses the cost allocation tags to organize your resource costs on your cost allocation report, to make it easier

for you to categorize and track your AWS costs. AWS provides two types of cost allocation tags, an AWS-generated tag and user-defined tags. AWS defines, creates, and applies the AWS-generated tag for you, and you define, create, and apply user-defined tags. You must activate both types of tags separately before they can appear in Cost Explorer or on a cost allocation report.

For more information on Cost Allocation tags, please visit the below URL: <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-allocotags.html>

NEW QUESTION 92

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

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

Answer: A

Explanation:

You can use Lambda functions as triggers for your Amazon DynamoDB table. Triggers are custom actions you take in response to updates made to the DynamoDB table. To create a trigger, first you enable Amazon DynamoDB Streams for your table. Then, you write a Lambda function to process the updates published to the stream.

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

NEW QUESTION 95

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

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

Answer: C

Explanation:

Answer - C

You can monitor the progress of a stack update by viewing the stack's events. The console's Cvents tab displays each major step in the creation and update of the stack sorted by the time of each event with latest events on top. The start of the stack update process is marked with an UPDATE_IN_PROGRESS event for the stack For more information on Monitoring your stack, please visit the below URL:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks-monitor-stack.html>

NEW QUESTION 99

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 103

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

- A. Create a global AWS CloudTrail Trail

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 104

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 107

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

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

Answer: B

Explanation:

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

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

- Your end users and IT administrators can use their existing corporate credentials to log on to AWS applications such as Amazon Workspaces, Amazon WorkDocs, or Amazon WorkMail.
- You can manage AWS resources like Amazon EC2 instances or Amazon S3 buckets through 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 109

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 110

Your CTO thinks your AWS account was hacked. What is the only way to know for certain if there was unauthorized access and what they did, assuming your hackers are very sophisticated AWS engineers and doing everything they can to cover their tracks?

- A. Use CloudTrail Log File Integrity Validation.
- B. Use AWS Config SNS Subscriptions and process events in real time.
- C. Use CloudTrail backed up to AWS S3 and Glacier.
- D. Use AWS Config Timeline forensics.

Answer: A

Explanation:

To determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it, you can use CloudTrail log file integrity validation. This feature is built using industry standard algorithms: SHA-256 for hashing and SHA-256 with RSA for digital signing. This makes it computationally infeasible to modify, delete or forge CloudTrail log files without detection. You can use the AWS CLI to validate the files in the location where CloudTrail delivered them

Validated log files are invaluable in security and forensic investigations. For example, a validated log file enables you to assert positively that the log file itself has not changed, or that particular user credentials performed specific API activity. The CloudTrail log file integrity validation process also lets you know if a log file has been deleted or changed, or assert positively that no log files were delivered to your account during a given period of time.

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

<http://docs.aws.amazon.com/awsccloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>

NEW QUESTION 112

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

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

Answer: B

Explanation:

One can use the CLI command list-access-keys to get the access keys. This command also returns the "CreateDate" of the keys. If the CreateDate is older than 2 months, then the keys can be deleted.

The Returns list-access-keys CLI command returns information about the access key IDs associated with the specified IAM user. If there are none, the action returns an empty list.

For more information on the CLI command, please refer to the below link: <http://docs.aws.amazon.com/cli/latest/reference/iam/list-access-keys.html>

NEW QUESTION 113

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 verify that there is an issue but will not help pinpoint the issue.

The Elastic Load Balancer logs will also just verify 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 115

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

- A. Terminate the existing Auto Scaling group
- B. Create a new launch configuration with the new Instance type
- C. Attach that to the new Autoscaling Group.

- D. Use the AutoScalingRollingUpdate policy on CloudFormation Template Auto Scalinggroup
- E. Use the Rolling Update feature which is available for EC2 Instances.
- F. Manually terminate the instances, launch new instances with the new instance type and attach them to the Autoscaling group

Answer: B

Explanation:

The AWS::AutoScaling::AutoScalingGroup resource supports an UpdatePolicy attribute. This is used to define how an Auto Scalinggroup resource is updated when an update to the Cloud Formation stack occurs. A common approach to updating an Auto Scaling group is to perform a rolling update, which is done by specifying the AutoScalingRollingUpdate policy. This retains the same Auto Scalinggroup and replaces old instances with new ones, according to the parameters specified. For more information on AutoScaling Rolling Update, please refer to the below link:

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

NEW QUESTION 116

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 119

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

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

Answer: C

Explanation:

The AWS Documentation mentions

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

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

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

NEW QUESTION 121

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 123

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

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

Answer: A

Explanation:

The AWS Documentation mentions

AWS Cloud Formation ensures all stack resources are created or deleted as appropriate. Because AWS CloudFormation treats the stack resources as a single unit, they must all be created or deleted successfully for the stack to be created or deleted. If a resource cannot be created, AWS CloudFormation rolls the stack back and automatically deletes any resources that were created. For more information on Cloudformation, please refer to the below link: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacks.html>

NEW QUESTION 124

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

- A. CodeCommit
- B. CodeDeploy
- C. S3Lifecycle
- D. Route53

Answer: B

Explanation:

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

NEW QUESTION 127

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

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

Answer: ABD

Explanation:

The AWS documentation mentions IAM supports AWS CodeCommit with three types of credentials: Git credentials, an IAM-generated user name and password pair you can use to communicate with AWS CodeCommit repositories over HTTPS. SSH keys, a locally generated public-private key pair that you can associate with your IAM user to communicate with AWS CodeCommit repositories over SSH. AWS access keys, which you can use with the credential helper included with the AWS CLI to communicate with AWS CodeCommit repositories over HTTPS. https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_ssh-keys.html

NEW QUESTION 128

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


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

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

Answer: C

Explanation:

The below diagram from the AWS documentation shows the deployment steps



For more information on the deployment steps please refer to the below link:
• <http://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-steps.html>

NEW QUESTION 130

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

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

Answer: AC

Explanation:

As per AWS documentation.

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

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

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

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

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

on the specified instances.

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

NEW QUESTION 134

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 136


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

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

Answer: B

Explanation:

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



The AWS Documentation mentions

- 1) With health-based rolling updates. Elastic Beanstalk waits until instances in a batch pass health checks before moving on to the next batch.
- 2) For time-based rolling updates, you can configure the amount of time that Elastic Beanstalk waits after completing the launch of a batch of instances before moving on to the next batch. This pause time allows your application to bootstrap and start serving requests.
- 3) Immutable environment updates are an alternative to rolling updates that ensure that configuration changes that require replacing instances are applied efficiently and safely. If an immutable environment update fails, the rollback process requires only terminating an Auto Scaling group. A failed rolling update, on the other hand, requires performing an additional rolling update to roll back the changes.

For more information on Rolling updates for Elastic beanstalk configuration updates, please visit the below URL:

- <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.ro11ingupdates.html>

NEW QUESTION 141

Which of the following environment types are available in the Elastic Beanstalk environment. Choose 2 answers from the options given below

- A. Single Instance
- B. Multi-Instance
- C. Load Balancing Autoscaling
- D. SQS, Autoscaling

Answer: AC

Explanation:

The AWS Documentation mentions
In Elastic Beanstalk, you can create a load-balancing, autoscaling environment or a single-instance environment. The type of environment that you require depends on the application that you deploy.
When you go onto the Configuration for your environment, you will be able to see the Environment type from there

NEW QUESTION 146

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 151

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


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

Answer: D

Explanation:

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

When you go to Cloudwatch events, you can use the Target as EC2 CreateSnapshot API call as shown below.
Create rules to invoke Targets based on Events happening in your AWS environment.



The AWS Documentation mentions
Amazon Cloud Watch Events delivers a near real-time stream of system events that describe changes in Amazon Web Services (AWS) resources. Using simple rules that you can quickly set up, you can match events and route them to one or more target functions or streams. Cloud Watch Events becomes aware of operational changes as they occur. Cloud Watch Events responds to these operational changes and takes corrective action as necessary, by sending messages to respond to the environment, activating functions, making changes, and capturing state information. For more information on Cloud watch Events, please visit the below URL:
• <http://docs.aws.amazon.com/AmazonCloudWatch/latest/events/WhatIsCloudWatchEvents.html>

NEW QUESTION 154

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

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

G. Create a user for the IT Auditor and ensure that full control is given to the user for the Cloudwatch logs.

Answer: B

Explanation:

The screenshot shows the AWS IAM console interface. On the left, under 'User details', the following information is displayed:

- User name: testpuser
- AWS access type: AWS Management Console access - with a password
- Console password type: Autogenerated
- Require password reset: Yes
- Permissions boundary: Permissions boundary is not set

On the right, under 'Permissions summary', it states: 'The following policies will be attached to the user shown above.'

Type	Name
Managed policy	AWSCloudTrailReadOnlyAccess
Managed policy	IAMUserChangePassword

There is a 'Learn More' button next to the 'Access' section.

The AWS Documentation clearly mentions the below

AWS CloudTrail is an AWS service that helps you enable governance, compliance, and operational and risk auditing of your AWS account. Actions taken by a user,

role, or an AWS service are recorded as events in CloudTrail. Events include actions taken in the AWS Management Console, AWS Command Line Interface, and AWS SDKs and APIs.

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

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

NEW QUESTION 157

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 161

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 162

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

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

Answer: A

Explanation:

The AWS Documentation mentions

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

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

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

NEW QUESTION 166

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

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

Answer: BC

Explanation:

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

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

The AWS Documentation mentions

When you launch an instance in Amazon EC2, you have the option of passing user data to the instance that can be used to perform common automated configuration tasks and even run scripts after the instance starts. You can pass two types of user data to Amazon EC2: shell scripts and cloud-init directives. You can

also pass this data into the launch wizard as plain text, as a file (this is useful for launching instances using the command line tools), or as base64-encoded text (for API calls).

For more information on User data please refer to the below link: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/user-data.html>

NEW QUESTION 170

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

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

Answer: D

Explanation:

Creating multiple cloudformation templates is an example of using nested stacks. The advantage of using nested stacks is given below as per the AWS documentation

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

other templates.

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

NEW QUESTION 172

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 175

Your company is planning to develop an application in which the front end is in .Net and the backend is in DynamoDB. There is an expectation of a high load on the application. How could you ensure the scalability of the application to reduce the load on the DynamoDB database? Choose an answer from the options below.

- A. Add more DynamoDB databases to handle the load.
- B. Increase write capacity of Dynamo DB to meet the peak loads
- C. Use SQS to assist and let the application pull messages and then perform the relevant operation in DynamoDB.
- D. Launch DynamoDB in Multi-AZ configuration with a global index to balance writes

Answer: C

Explanation:

When the idea comes for scalability then SQS is the best option. Normally DynamoDB is scalable, but since one is looking for a cost effective solution, the messaging in SQS can assist in managing the situation mentioned in the question.

Amazon Simple Queue Service (SQS) is a fully-managed message queuing service for reliably communicating among distributed software components and microservices - at any scale. Building applications from individual components that each perform a discrete function improves scalability and reliability, and is best practice design for modern applications. SQS makes it simple and cost-effective to decouple and coordinate the components of a cloud application. Using SQS, you can send, store, and receive messages between software components at any volume, without losing messages or requiring other services to be always available

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

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

NEW QUESTION 179

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

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

Answer: B

Explanation:

The AWS documentation mentions

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

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

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

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

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

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

NEW QUESTION 183

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

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

Answer: A

Explanation:

The AWS documentation mentions the following about Opswork stack

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

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

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

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

NEW QUESTION 186

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

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

Answer: D

Explanation:

If you refer to the documentation, you will see that Resources is the only mandatory field Specifies the stack resources and their properties, such as an Amazon Elastic Compute Cloud instance or an Amazon Simple Storage Service bucket. For more information on cloudformation templates, please refer to the below link:
• <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-anatomy.html>

NEW QUESTION 188

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 189

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

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

Answer: B

Explanation:

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

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

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

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

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

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

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

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

NEW QUESTION 193

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 198

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

- A. It is not possible that the stack creation will wait until one service is created and launched.
- B. The user can use the HoldCondition resource to wait for the creation of the other dependent resources.
- C. The user can use the WaitCondition resource to hold the creation of the other dependent resources.
- D. The user can use the DependentCondition resource to hold the creation of the other dependent resources.

Answer: D

Explanation:

You can use a wait condition for situations like the following:

To coordinate stack resource creation with configuration actions that are external to the stack creation

To track the status of a configuration process

For more information on CloudFormation Wait condition please visit the link

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-properties-waitcondition.html>

NEW QUESTION 200

You are writing an AWS Cloud Formation template and you want to assign values to properties that will not be available until runtime. You know that you can use intrinsic functions to do this but are unsure as to which part of the template they can be used in. Which of the following is correct in describing how you can currently use intrinsic functions in an AWS CloudFormation template?

- A. You can use intrinsic functions in any part of a template.
- B. You can only use intrinsic functions in specific parts of a template.
- C. You can use intrinsic functions in resource properties, metadata attributes, and update policy attributes.
- D. You can use intrinsic functions only in the resource properties part of a template.
- E. You can use intrinsic functions in any part of a template, except AWSTemplateFormatVersion and Description.

Answer: B

Explanation:

This is clearly given in the AWS documentation. Intrinsic Function Reference

AWS Cloud Formation provides several built-in functions that help you manage your stacks. Use intrinsic functions in your templates to assign values to properties that are not available until runtime. Note

You can use intrinsic functions only in specific parts of a template. Currently, you can use intrinsic functions in resource properties, outputs, metadata attributes, and update policy attributes. You can also use intrinsic functions to conditionally create stack resources. For more information on intrinsic function please refer to the below link <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference.html>

NEW QUESTION 205

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 209

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

Productionaccount.

- B. When creating the role, define the Development account as a trustedentity and specify a permissions policy that allows trusted users to update theS3 bucket.
- C. Use web identity federation with a third-partyidentity provider with AWS STS to grant temporary credentials and membershipinto the production 1AM user.
- D. Createan 1AM cross account role in the Production account that allows users from theDevelopment 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 213

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your DOP-C01 Exam with Our Prep Materials Via below:

<https://www.certleader.com/DOP-C01-dumps.html>