

## DOP-C01 Dumps

### AWS Certified DevOps Engineer- Professional

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



**NEW QUESTION 1**

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

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

**Answer: A**

**Explanation:**

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

**NEW QUESTION 2**

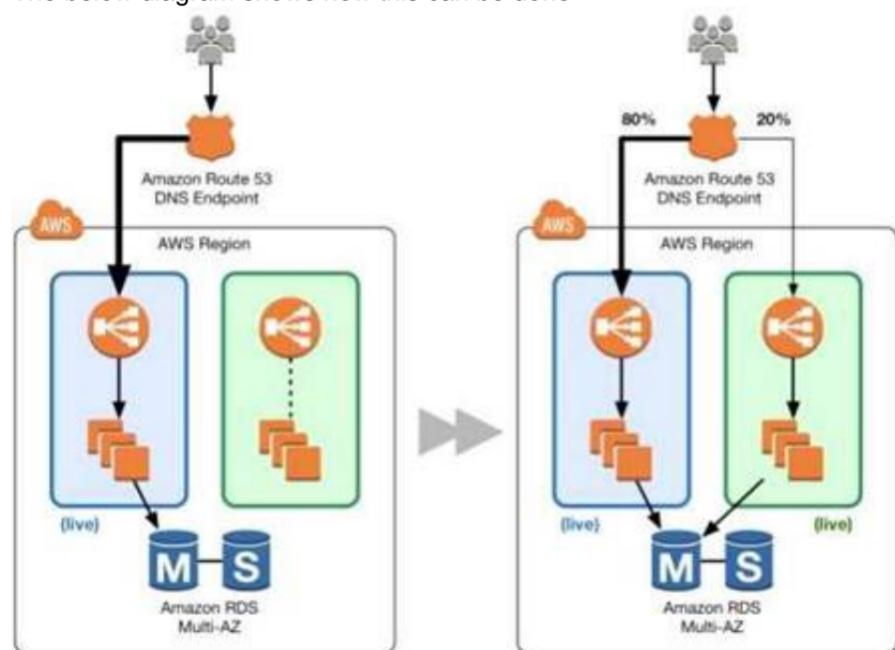
Your application is currently running on Amazon EC2 instances behind a load balancer. Your management has decided to use a Blue/Green deployment strategy. How should you implement this for each deployment?

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

**Answer: C**

**Explanation:**

The below diagram shows how this can be done



- 1) First create a new ELB which will be used to point to the new production changes.
- 2) Use the Weighted Route policy for Route53 to distribute the traffic to the 2 ELB's based on a 80- 20% traffic scenario. This is the normal case, the % can be changed based on the requirement.
- 3) Finally when all changes have been tested, Route53 can be set to 100% for the new ELB.

Option A is incorrect because this is a failover scenario and cannot be used for Blue green deployments. In Blue Green deployments, you need to have 2 environments running side by side. Option B is incorrect, because you need to have a production stack with the changes which will run side by side. Option D is incorrect because this is not a blue green deployment scenario. You cannot control which users will go the new EC2 instances.

For more information on blue green deployments, please refer to the below document link: from AWS [https://dOawsstatic.com/whitepapers/AWS\\_Blue\\_Green\\_Deployments.pdf](https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf)

**NEW QUESTION 3**

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

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

**Answer: D**

**Explanation:**

If you have custom health checks, you can send the information from your health checks to Auto Scaling so that Auto Scaling can use this information. For

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

**NEW QUESTION 4**

You have just recently deployed an application on EC2 instances behind an ELB. After a couple of weeks, customers are complaining on receiving errors from the application. You want to diagnose the errors and are trying to get errors from the ELB access logs. But the ELB access logs are empty. What is the reason for this.

- A. You do not have the appropriate permissions to access the logs
- B. You do not have your CloudWatch metrics correctly configured
- C. ELB Access logs are only available for a maximum of one week.
- D. Access logging is an optional feature of Elastic Load Balancing that is disabled by default

**Answer: D**

**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.

Access logging is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logging for your load balancer, Elastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify. You can disable access logging at any time.

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 5**

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/launch\\_configuration.html](http://docs.aws.amazon.com/autoscaling/latest/userguide/launch_configuration.html)

**NEW QUESTION 6**

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

**NEW QUESTION 7**

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

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

**Answer: C**

**Explanation:**

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

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

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

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

Option D is invalid because there is no Auto Scaling template.

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

**NEW QUESTION 8**

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

**Answer: DG**

**Explanation:**

The scaling of EC2 instances in the Auto Scaling group is normally done with the metric of the CPU utilization of the current instances in the Auto Scaling group

For more information on scaling in your Auto Scaling Group, please refer to the below link:

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

Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) instances, making them a natural fit for production database workloads. When you provision a Multi-AZ DB instance, Amazon RDS automatically creates a primary DB instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure, Amazon RDS performs an automatic failover to the standby (or to a read replica in the case of Amazon Aurora), so that you can resume database operations as soon as the failover is complete. For more information on RDS Multi-AZ please refer to the below link:

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

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

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

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

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

Option F is invalid because increasing Read-Replicas will not suffice for write operations if the primary DB fails.

**NEW QUESTION 9**

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

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

**Answer: C**

**Explanation:**

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

Option A and B are invalid because Autoscaling policies are not designed for these purposes. Option D is invalid because you use Cloudwatch Logs Agent and not the monitoring agent. For more information on Cloudwatch logs, please refer to the below link: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html>

**NEW QUESTION 10**

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

- A. Add an Amazon S3 bucket policy with a condition statement to allow access only from Amazon EC2 instances with RFC 1918 IP addresses and enable bucket versioning.
- B. Add an Amazon S3 bucket policy with a condition statement that requires multi-factor authentication in order to delete objects and enable bucket versioning.
- C. Use a configuration management service to deploy AWS Identity and Access Management user credentials to the Amazon EC2 instance
- D. Use these credentials to securely access the Amazon S3 bucket when deploying code.
- E. Create an Amazon Identity and Access Management role with authorization to access the Amazon S3 bucket, and launch all of your application's Amazon EC2 instances with this role.
- F. Use AWS Data Pipeline to lifecycle the data in your Amazon S3 bucket to Amazon Glacier on a weekly basis.
- G. Use AWS Data Pipeline with multi-factor authentication to securely deploy code from the Amazon S3 bucket to your Amazon EC2 instances.

**Answer:** BD

**Explanation:**

You can add another layer of protection by enabling MFA Delete on a versioned bucket. Once you do so, you must provide your AWS account's access keys and a valid code from the account's MFA device in order to permanently delete an object version or suspend or reactivate versioning on the bucket. For more information on MFA please refer to the below link: <https://aws.amazon.com/blogs/security/securing-access-to-aws-using-mfa-part-3/>

IAM roles are designed so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that the applications use. Instead of creating and distributing your AWS credentials, you can delegate permission to make API requests using IAM roles. For more information on Roles for EC2 please refer to the below link: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

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

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

**NEW QUESTION 10**

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

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

**Answer:** CE

**Explanation:**

The AWS Security Token Service (STS) is a web service that enables you to request temporary, limited-privilege credentials for AWS Identity and Access Management (IAM) users or for users that you authenticate (federated users). The token can then be used to grant access to the objects in S3. You can then provide access to the objects based on the key values generated via the user id. Option A is possible but then becomes a maintenance overhead because of the number of buckets. Option B is invalid because IAM users is not a good security practice.

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

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

**NEW QUESTION 11**

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

Amazon OastlCache is an in-memory key/value store that sits between ycbetappiicipiJGra arcdalie data store (database) that it accesses. Whenever your application requests data, it first makes the request to the DastlCache cache. If the data exists in the cache and is current, OastlCache 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 Elastic Cache 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 16

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

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

**Answer: B**

#### Explanation:

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

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

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

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

- <https://jenkins.io/>

#### NEW QUESTION 17

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. Dockerrun.aws.json
- B. .ebextensions
- C. Dockerrun.json
- D. Dockerfile

**Answer: A**

#### Explanation:

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

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

the containers to mount. [http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create\\_deploy\\_docker\\_v2config.html](http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_v2config.html)

#### NEW QUESTION 21

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

- A. Publish your data to Cloud Watch Logs, and configure your application to autoscale to handle the load on demand.
- B. Publish your log data to an Amazon S3 bucket
- C. Use AWS CloudFormation to create an Auto Scalinggroup to scale your post-processing application which is configured to pull down your log files stored an Amazon S3.
- D. Post your log data to an Amazon Kinesis data stream, and subscribe your log-processing application so that is configured to process your logging data.
- E. Create a multi-AZ Amazon RDS MySQL cluster, post the logging data to MySQL, and run a map reduce job to retrieve the required information on user counts.

**Answer: C**

#### Explanation:

When you see Amazon Kinesis as an option, this becomes the ideal option to process data in real time.

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

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

below URL:

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

#### NEW QUESTION 24

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 CC2 creates snapshots of your instance's root volume and any other CBS volumes attached to your instance. New CBS 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 CBS initialization please refer to the below link:

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

#### NEW QUESTION 26

You have a complex system that involves networking, IAM policies, and multiple, three-tier applications. You are still receiving requirements for the new system, so you don't yet know how many AWS components will be present in the final design. You want to start using AWS CloudFormation to define these AWS resources so that you can automate and version-control your infrastructure. How would you use AWS CloudFormation to provide agile new environments for your customers in a cost-effective, reliable manner?

- A. Manually create one template to encompass all the resources that you need for the system, so you only have a single template to version-control.
- B. Create multiple separate templates for each logical part of the system, create nested stacks in AWS CloudFormation, and maintain several templates to version-control.
- C. ➤/
- D. Create multiple separate templates for each logical part of the system, and provide the outputs from one to the next using an Amazon Elastic Compute Cloud (EC2) instance running the SDK for finer granularity of control.
- E. Manually construct the networking layer using Amazon Virtual Private Cloud (VPC) because this does not change often, and then use AWS CloudFormation to define all other ephemeral resources.

**Answer: B**

#### Explanation:

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

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

#### NEW QUESTION 27

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

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

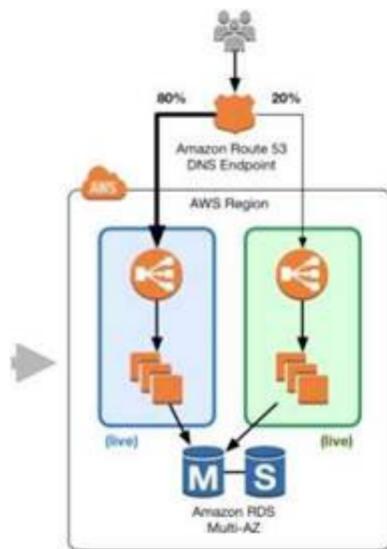
Which combination of services would give you the ability to control traffic between different deployed versions of your application?

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

**Answer: B**

#### Explanation:

This is an example of Blue green deployment.



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

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

- [https://dOawsstatic.com/whitepapers/AWS\\_Blue\\_Green\\_Deployments.pdf](https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf)

### NEW QUESTION 29

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?

- Use CloudFormation templates
- Configuration Management Templates
- Saved Configurations
- 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 33

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

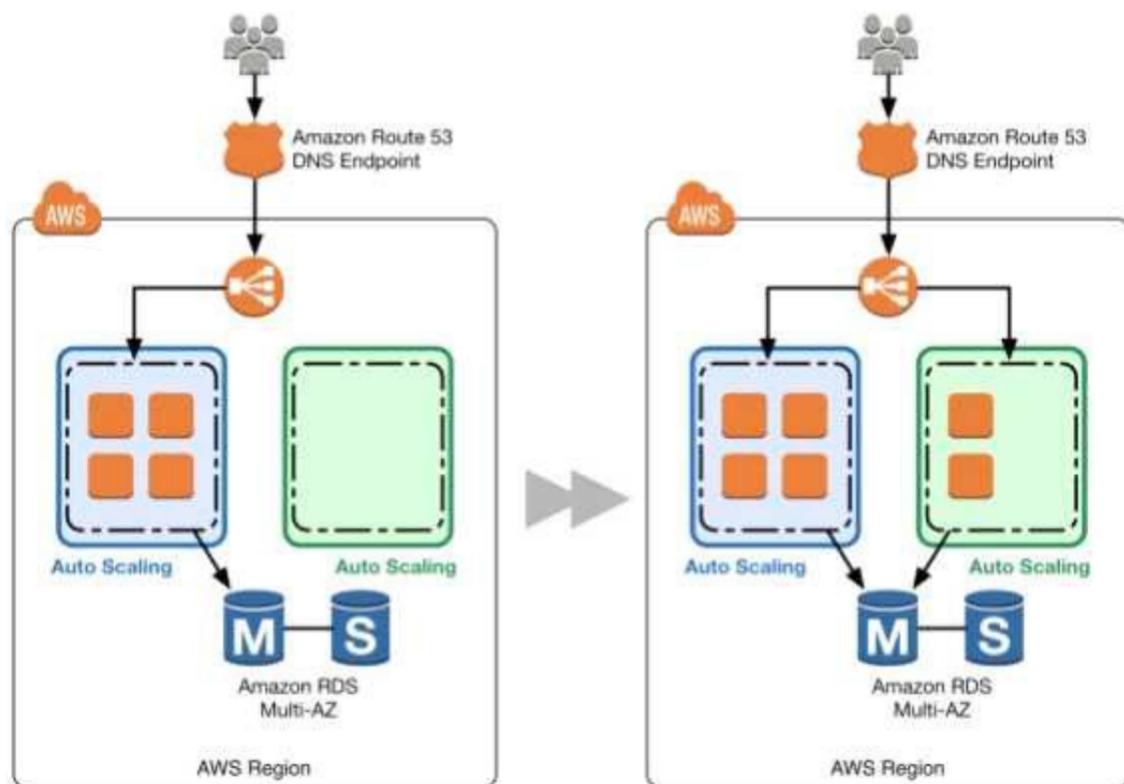
- Use Newest Instance to phase out all instances that use the previous configuration.
- Attach an additional ELB to your Auto Scaling configuration and phase in newer instances while removing older instances.
- Use OldestLaunchConfiguration to phase out all instances that use the previous configuration.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

**Answer: CD**

#### Explanation:

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

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



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

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

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

- [https://dOawsstatic.com/whitepapers/AWS\\_Blue\\_Green\\_Deployments.pdf](https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf)

#### NEW QUESTION 34

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

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

**Answer: B**

#### Explanation:

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

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

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

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

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

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

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

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

#### NEW QUESTION 35

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?

- Develop models of your entire cloud system in CloudFormation
- Use this model in Staging and Production to achieve greater parity
- \*/
- Use AWS Config to force the Staging and Production stacks to have configuration parity
- Any differences will be detected for you so you are aware of risks.
- 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.
- Use AWS ECS and Docker clustering
- 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 40**

What is web identity federation?

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

**Answer: B**

**Explanation:**

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

[http://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles\\_providers\\_oidc.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html)

**NEW QUESTION 45**

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 49**

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

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

**Answer: C**

**Explanation:**

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

[https://media.amazonwebservices.com/architecturecenter/AWS\\_ac\\_ra\\_ftha\\_04.pdf](https://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_ftha_04.pdf)

**NEW QUESTION 53**

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 58

You are creating a new API for video game scores. Reads are 100 times more common than writes, and the top 1% of scores are read 100 times more frequently than the rest of the scores. What's the best design for this system, using DynamoDB?

- A. DynamoDB table with 100x higher read than write throughput, with CloudFront caching.
- B. DynamoDB table with roughly equal read and write throughput, with CloudFront caching.
- C. DynamoDB table with 100x higher read than write throughput, with ElastiCache caching.
- D. DynamoDB table with roughly equal read and write throughput, with ElastiCache caching.

**Answer: D**

#### Explanation:

Because the IOOx read ratio is mostly driven by a small subset, with caching, only a roughly equal number of reads to writes will miss the cache, since the supermajority will hit the top 1% scores. Knowing we need to set the values roughly equal when using caching, we select AWS OastiCache, because CloudFront cannot directly cache DynamoDB queries, and OastiCache is an excellent in-memory cache for database queries, rather than a distributed proxy cache for content delivery.

For more information on DynamoDB table gudelines please refer to the below link:

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

#### NEW QUESTION 60

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

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

**Answer: A**

#### Explanation:

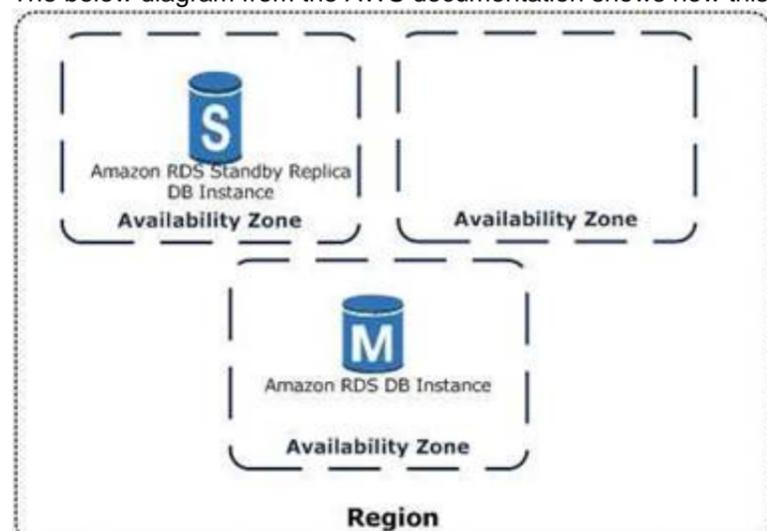
Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB)

Instances, making them a natural fit for production database

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

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

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



Option B is invalid because the replication is synchronous.

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

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

#### NEW QUESTION 64

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 68**

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\_CLEANUPJN\_PROGRESS
- B. DELETE\_COMPLETE
- C. ROLLBACK\_IN\_PROGRESS
- D. UPDATE\_IN\_PROGRESS

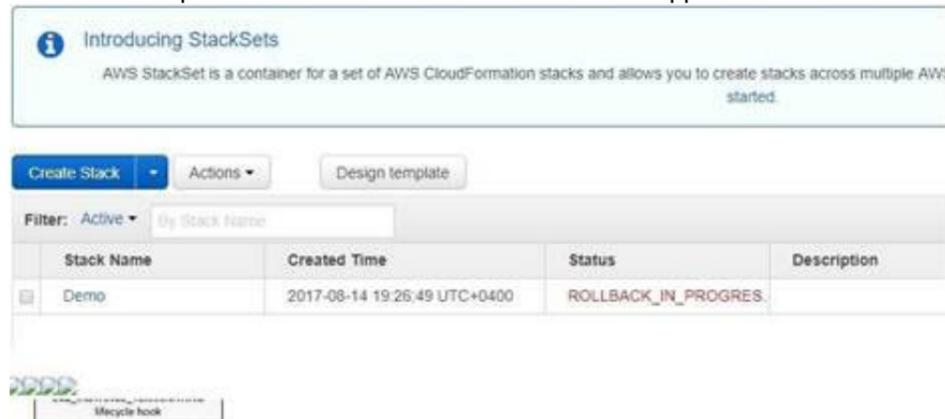
**Answer: C**

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

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

**Answer: B**

**Explanation:**

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

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

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

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

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

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

**NEW QUESTION 75**

You are creating an application which stores extremely sensitive financial information. All information in the system must be encrypted at rest and in transit. Which

of these is a violation of this policy?

- A. ELB SSL termination.
- B. ELB Using Proxy Protocol v1.
- C. CloudFront Viewer Protocol Policy set to HTTPS redirection.
- D. Telling S3 to use AES256 on the server-side.

**Answer:** A

**Explanation:**

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

**NEW QUESTION 79**

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

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

Which services would suffice the above requirements

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

**Answer:** C

**Explanation:**

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

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

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

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

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

- [http://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_credentials\\_getting-report.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_getting-report.html)

**NEW QUESTION 83**

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 86**

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 later.
- G. Create a Memcached cluster in AWS ElastiCache.
- H. Create cache logic to serve requests which can be served later 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 88**

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

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

**Answer: B**

**Explanation:**

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

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

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

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

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

**NEW QUESTION 92**

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

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

**Answer: B**

**Explanation:**

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

These changes can range from simple configuration changes, such

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

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

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

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/updates-through.html>

**NEW QUESTION 97**

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

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

**Answer: B**

**Explanation:**

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

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

visit the below URL:

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

**NEW QUESTION 101**

You need your API backed by DynamoDB to stay online during a total regional AWS failure. You can tolerate a couple minutes of lag or slowness during a large failure event, but the system should recover with normal operation after those few minutes. What is a good approach?

- A. Set up DynamoDB cross-region replication in a master-standby configuration, with a single standby in another region
- B. Create an Auto Scaling Group behind an ELB in each of the two regions for your application layer in which DynamoDB is running
- C. Add a Route53 Latency DNS Record with DNS Failover, using the ELBs in the two regions as the resource records.
- D. Set up a DynamoDB Global table
- E. Create an Auto Scaling Group behind an ELB in each of the two regions for your application layer in which the DynamoDB is running
- F. Add a Route53 Latency DNS Record with DNS Failover, using the ELBs in the two regions as the resource records.
- G. Set up a DynamoDB Multi-Region table
- H. Create a cross-region ELB pointing to a cross-region Auto Scaling Group, and direct a Route53 Latency DNS Record with DNS Failover to the cross-region ELB.
- I. Set up DynamoDB cross-region replication in a master-standby configuration, with a single standby in another region
- J. Create a cross-region ELB pointing to a cross-region Auto Scaling Group, and direct a Route53 Latency DNS Record with DNS Failover to the cross-region ELB.

**Answer: B**

**Explanation:**

Updated based on latest AWS updates

Option A is invalid because using Latency based routing will send traffic on the region with the standby instance. This is an active/passive replication and you can't write to the standby table unless there is a failover. Answer A can work only if you use a failover routing policy.

Option D is invalid because there is no concept of a cross-region CLB.

Amazon DynamoDB Global Tables provide a fully managed solution for deploying a multi-region, multi-master database, without having to build and maintain your own replication solution. When you create a global table, you specify the AWS regions where you want the table to be available. DynamoDB performs all of the necessary tasks to create identical tables in these regions, and propagate ongoing data changes to all of them.

For more information on DynamoDB Global Tables, please visit the below URL:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GlobalTables.html>

**NEW QUESTION 102**

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

- A. Use a long-poll on the Resources object in your CloudFormation stack and display those state changes in the UI for the system.
- B. Use a long-poll on the ListStacks API 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 Events 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 107**

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

- A. Build the application out using AWS Cognito and web identity federation to allow users to log in using Facebook or Google Account
- B. Once they are logged in, the secret token passed to that user is used to directly access resources on AWS, like AWS S3.
- C. Use JWT or SAML compliant systems to build authorization policies
- D. Users log in with a username and password, and are given a token they can use indefinitely to make calls against the photo infrastructure.
- E. Use AWS API Gateway with a constantly rotating API Key to allow access from the client-side
- F. Construct a custom build of the SDK and include S3 access in it.
- G. Create an AWS OAuth Service Domain and grant public sign-up and access to the domain
- H. 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 109**

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

- A. AWS CloudFormation
- B. AWS OpsWorks

- C. AWS ELB+ EC2 with CLI Push
- D. AWS Elastic Beanstalk

**Answer:** D

**Explanation:**

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

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

**NEW QUESTION 112**

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

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

**Answer:** D

**Explanation:**

Custom resources enable you to write custom provisioning logic in templates that AWS Cloud Formation runs anytime you create, update (if you changed the custom resource), or delete stacks. For example, you might want to include resources that aren't available as AWS Cloud Formation resource types. You can include those resources by using custom resources. That way you can still manage all your related resources in a single stack. Use the AWS::CloudFormation::CustomResource or Custom::String resource type to define custom resources in your templates. Custom resources require one property: the service token, which specifies where AWS CloudFormation sends requests to, such as an Amazon SNS topic. For more information on Custom Resources in CloudFormation, please visit the below URL:

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

**NEW QUESTION 116**

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 121**

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 Scalinggroup
- B. Create a new launch configuration with the new Instance typ
- C. Attach that to the new Autoscaing 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 125**

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

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

**Answer: B**

**Explanation:**

The AWS Documentation mentions

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

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

**NEW QUESTION 127**

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

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

**Answer: ABD**

**Explanation:**

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

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

**NEW QUESTION 131**

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 132**

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 136**

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

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

**Answer:**

AB

**Explanation:**

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

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

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

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

**NEW QUESTION 137**

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 140**

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 145**

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

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

**Answer: B****Explanation:**

The AWS Documentation mentions

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

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

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

**NEW QUESTION 150**

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 152**

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

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

**Answer: B**

**Explanation:**

The AWS Documentation mentions

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

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

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

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

**NEW QUESTION 157**

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 158**

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

Desired capacity - 2 Maximum capacity - 4

The current number of instances running in the Autoscaling Group is 2. You have been notified that for a duration of an hour, you need to ensure that no new instances are launched by the Autoscaling Group Which of the below 2 actions can be carried out to fulfil this requirement

- A. Change the Maximum capacity to 2
- B. Change the Desired capacity to 4
- C. Suspend the Launch process of the Autoscaling Group
- D. Change the Minimum capacity to 2

**Answer: AC**

**Explanation:**

You can temporarily suspend the creation of new instances by either reducing the Maximum capacity to 2, so that the current instances running which is 2, matches the maximum limit.

Secondly you can suspend the launch process of the Autoscaling Group The AWS Documentation mentions Scaling Processes

Amazon EC2 Auto Scaling supports the following scaling processes:

Launch Adds a new EC2 instance to the group, increasing its capacity. Warning

If you suspend Launch, this disrupts other processes. For example, you can't return an instance in a standby state to service if the Launch process is suspended, because the group can't scale. For more information on Autoscaling Suspend and Resume, please visit the below URL:

- ? <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resume-processes.html>

**NEW QUESTION 163**

Your company has a set of EC2 resources hosted on AWS. Your new IT procedures state that AWS EC2 Instances must be of a particular Instance type. Which of the following can be used to get the list of EC2 Instances which currently don't match the instance type specified in the new IT procedures

- A. Use AWS Cloudwatch alarms to check which EC2 Instances don't match the intended instance type.
- B. Use AWS Config to create a rule to check the EC2 Instance type
- C. Use Trusted Advisor to check which EC2 Instances don't match the intended instance type.
- D. Use VPC Flow Logs to check which EC2 Instances don't match the intended instance type.

**Answer: B**

**Explanation:**

In AWS Config, you can create a rule which can be used to check if EC2 instances follow a particular instance type. Below is a snapshot of the output of a rule to check if EC2 instances matches the type of t2micro.

**Description** Checks whether your EC2 instances are of the specified instance types.

**Trigger type** Configuration changes

**Scope of changes** Resources

**Resource types** EC2 Instance

**Config rule ARN** arn:aws:config:ap-southeast-1:213171387512:config-rule/config-rule-1exnif

**Parameters** instanceType: t2.micro

**Overall rule status** Last successful invocation on August 29, 2017 at 12:06:57 PM

Last successful evaluation on August 29, 2017 at 12:06:59 PM

**Resources evaluated**

Click on the icon to view configuration details for the resource when it was last evaluated with this rule.

Resource type	Config timeline	Compliance	Last successful invocation	Last successful evaluation	Manage resource
EC2 Instance	i-0798cb2c515b3aa50	Noncompliant	August 29, 2017 11:25:10 AM	August 29, 2017 11:25:12 AM	
EC2 Instance	i-0f2a1120bec8db3d0	Noncompliant	August 29, 2017 12:06:57 PM	August 29, 2017 12:06:59 PM	

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

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

**NEW QUESTION 168**

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

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

**Answer: AD**

**Explanation:**

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

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

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

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

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

**NEW QUESTION 171**

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

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

**Answer: C**

**Explanation:**

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

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

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

**NEW QUESTION 174**

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 178**

An enterprise wants to use a third-party SaaS application running on AWS.. The SaaS application needs to have access to issue several API commands to discover Amazon EC2 resources running within the enterprise's account. The enterprise has internal security policies that require any outside access to their environment must conform to the principles of least privilege and there must be controls in place to ensure that the credentials used by the SaaS vendor cannot be used by any other third party. Which of the following would meet all of these conditions?

- A. From the AWS Management Console, navigate to the Security Credentials page and retrieve the access and secret key for your account.
- B. Create an IAM user within the enterprise account assign a user policy to the IAM user that allows only the actions required by the SaaS applicatio
- C. Create a new access and secret key for the user and provide these credentials to the SaaS provider.
- D. Create an IAM role for cross-account access allows the SaaS provider's account to assume the role and assign it a policy that allows only the actions required by the SaaS application.
- E. Create an IAM role for EC2 instances, assign it a policy that allows only the actions required for the SaaS application to work, provide the role ARN to the SaaS provider to use when launching their application instances.

**Answer:** C

**Explanation:**

Many SaaS platforms can access aws resources via a Cross account access created in aws. If you go to Roles in your identity management, you will see the ability to add a cross account role.

Select Role Type



For more information on cross account role, please visit the below URL:

- [http://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial\\_cross-account-with-roles.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html)

**NEW QUESTION 182**

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 OastiCache 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 OastiCache, please visit the below link:

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

**NEW QUESTION 183**

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 186**

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

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

**Answer: B**

**Explanation:**

The AWS documentation mentions You can monitor your AWS costs by using Cloud Watch. With Cloud Watch, you can create billing alerts that notify you when your usage of your services exceeds thresholds that you define. You specify these threshold amounts when you create the billing alerts. When your usage exceeds these amounts, AWS sends you an email notification. You can also sign up to receive notifications when AWS prices change. For more information on billing alarms, please refer to the below URL:  
• <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/monitor-charges.html>

**NEW QUESTION 188**

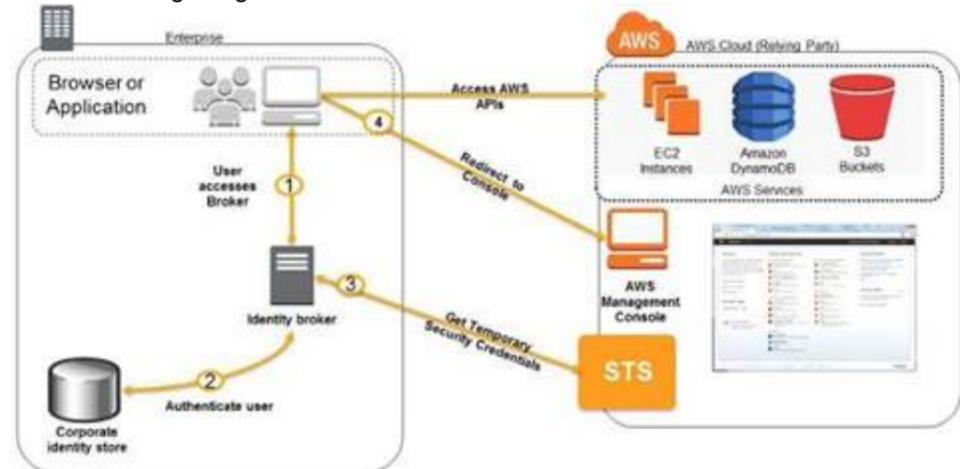
Which of the following will you need to consider so you can set up a solution that incorporates single sign-on from your corporate AD or LDAP directory and restricts access for each user to a designated user folder in a bucket? Choose 3 Answers from the options below

- A. Setting up a federation proxy or identity provider
- B. Using AWS Security Token Service to generate temporary tokens
- C. Tagging each folder in the bucket
- D. Configuring 1AM role
- E. Setting up a matching 1AM user for every user in your corporate directory that needs access to a folder in the bucket

**Answer: ABD**

**Explanation:**

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



For more information on federated access, please visit the below link: [http://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles\\_com\\_mon-scenarios\\_federated-users.htm](http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_com_mon-scenarios_federated-users.htm)

[https://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles\\_create\\_for-idp\\_saml.html?icmpid=docs\\_iam\\_console](https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_create_for-idp_saml.html?icmpid=docs_iam_console)

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

**NEW QUESTION 192**

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

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

**Answer:** C

**Explanation:**

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

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

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

**NEW QUESTION 195**

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

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

**Answer:** AB

**Explanation:**

The AWS Documentation mentions the following on AWS Pipeline

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

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

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

**NEW QUESTION 200**

You were just hired as a DevOps Engineer for a startup. Your startup uses AWS for 100% of their infrastructure. They currently have no automation at all for deployment, and they have had many failures while trying to deploy to production. The company has told you deployment process risk mitigation is the most important thing now, and you have a lot of budget for tools and AWS resources.

Their stack includes a 2-tier API with data stored in DynamoDB or S3, depending on type. The Compute layer is EC2 in Auto Scaling Groups. They use Route53 for DNS pointing to an ELB. An ELB balances load across the EC2 instances. The scaling group properly varies between 4 and 12 EC2 servers. Which of the following approaches, given this company's stack and their priorities, best meets the company's needs?

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

**Answer:** B

**Explanation:**

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

The AWS Documentation mentions the below on nested stacks

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

For more information on CloudFormation best practices, please visit the link:

- <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html> For more information on Blue Green Deployment, please visit the link:
- [https://dOawsstatic.com/whitepapers/AWS\\_Blue\\_Green\\_Deployments.pdf](https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf)

**NEW QUESTION 204**

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

- A. Create an OpsWorks Stack with two Layer
- B. The first contains lifecycle scripts for launching and bootstrapping an HTTP API on G2 instances for image processing, and the second has an always-on instance which monitors the S3 manifest bucket for new file
- C. When a new file is detected, request instances to boot on the artificial neural network layer
- D. When the instances are booted and the HTTP APIs are up, submit processing requests to individual instances.
- E. Make an S3 notification configuration which publishes to AWS Lambda on the manifest bucket
- F. Make the Lambda create a CloudFormation Stack which contains the logic to construct an autoscaling worker tier of EC2 G2 instances with the artificial neural network code on each instance
- G. Create an SQS queue of the images in the manifest
- H. Tear the stack down when the queue is empty.
- I. Deploy your artificial neural network code to AWS Lambda as a bundled binary for the C++ extension
- J. Make an S3 notification configuration on the manifest, which publishes to another AWS Lambda running controller code
- K. This controller code publishes all the images in the manifest to AWS Kinesis

- L. Your ANN code Lambda Function uses the Kinesis as an Event Source
- M. The system automatically scales when the stream contains image events.
- N. Create an Auto Scaling Load Balanced Elastic Beanstalk worker tier Application and Environment
- O. Deploy the artificial neural network code to G2 instances in this tier
- P. Set the desired capacity to 1. Make the code periodically check S3 for new manifests
- Q. When a new manifest is detected, push all of the images in the manifest into the SQS queue associated with the Elastic Beanstalk worker tier.

**Answer: B**

**Explanation:**

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

The AWS Documentation mentions the following on AWS S3 and Lambda

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

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

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

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

**NEW QUESTION 209**

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

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

**Answer: B**

**Explanation:**

The AWS documentation mentions the following

You can manage Amazon SNS 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 SNS messages with Amazon S3, use the Amazon SNS 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 SNS and sending larger messages please visit the link

**NEW QUESTION 210**

You are writing an AWS CloudFormation 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 CloudFormation 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 functions please refer to the below link <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference.html>

**NEW QUESTION 211**

Which of the following run command types are available for OpsWorks stacks? Choose 3 answers from the options given below.

- A. UpdateCustom Cookbooks
- B. Execute Recipes
- C. Configure
- D. UnDeploy

**Answer: ABC**

**NEW QUESTION 214**

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

- A. Create separate dashboards in every region

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

**Answer: C**

**Explanation:**

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

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

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

[http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/cross\\_region\\_dashboard.html](http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/cross_region_dashboard.html)

**NEW QUESTION 216**

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 221**

A vendor needs access to your AWS account. They need to be able to read protected messages in a private S3 bucket. They have a separate AWS account. Which of the solutions below is the best way to do this?

- A. Allow the vendor to ssh into your EC2 instance and grant them an IAM role with full access to the bucket.
- B. Create a cross-account IAM role with permission to access the bucket, and grant permission to use the role to the vendor AWS account.
- C. Create an IAM User with API Access Key
- D. Give the vendor the AWS Access Key ID and AWS Secret Access Key for the user.
- E. Create an S3 bucket policy that allows the vendor to read from the bucket from their AWS account.

**Answer: B**

**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](http://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html) | <https://docs.aws.amazon.com/AmazonS3/latest/dev/example-walkthroughs-managing-access-example2.html>

**NEW QUESTION 224**

You need to store a large volume of data. The data needs to be readily accessible for a short period, but then needs to be archived indefinitely after that. What is a cost-effective solution?

- A. Store all the data in S3 so that it can be more cost effective
- B. Store your data in Amazon S3, and use lifecycle policies to archive to Amazon Glacier
- C. Store your data in an EBS volume, and use lifecycle policies to archive to Amazon Glacier.
- D. Store your data in Amazon S3, and use lifecycle policies to archive to S3-Infrequently Access

**Answer: B**

**Explanation:**

The AWS documentation mentions the following on Lifecycle policies

Lifecycle configuration enables you to specify the lifecycle management of objects in a bucket. The configuration is a set of one or more rules, where each rule defines an action for Amazon S3 to apply to a group of objects. These actions can be classified as follows:

Transition actions - In which you define when objects transition to another storage class. For example, you may choose to transition objects to the STANDARD\_IA (IA, for infrequent access) storage class 30 days after creation, or archive objects to the GLACIER storage class one year after creation.

Expiration actions - In which you specify when the objects expire. Then Amazon S3 deletes the expired objects on your behalf. For more information on S3

Lifecycle policies, please visit the below URL

• <http://docs.aws.amazon.com/AmazonS3/latest/dev/object-lifecycle-mgmt.html>

**NEW QUESTION 227**

.....

## 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>