



# Amazon-Web-Services

## Exam Questions DOP-C01

AWS Certified DevOps Engineer- Professional

#### 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

You currently have the following setup in AWS

- 1) An Elastic Load Balancer
- 2) Auto Scaling Group which launches EC2 Instances
- 3) AMIs with your code pre-installed

You want to deploy the updates of your app to only a certain number of users. You want to have a cost-effective solution. You should also be able to revert back quickly. Which of the below solutions is the most feasible one?

- A. Create a second ELB, and a new Auto Scaling Group assigned a new Launch Configuration
- B. Create a new AMI with the updated app
- C. Use Route53 Weighted Round Robin records to adjust the proportion of traffic hitting the two ELBs.
- D. Create new AMIs with the new app
- E. Then use the new EC2 instances in half proportion to the older instances.
- F. Redeploy with AWS Elastic Beanstalk and Elastic Beanstalk version
- G. Use Route 53 Weighted Round Robin records to adjust the proportion of traffic hitting the two ELBs
- H. Create a full second stack of instances, cut the DNS over to the new stack of instances, and change the DNS back if a rollback is needed.

**Answer:** A

#### Explanation:

The Weighted Routing policy of Route53 can be used to direct a proportion of traffic to your application. The best option is to create a second CLB, attach the new Autoscaling Group and then use Route53 to divert the traffic.

Option B is wrong because just having EC2 instances running with the new code will not help.

Option C is wrong because Elastic beanstalk is good for development environments, and also there is no mention of having 2 environments where environment url's can be swapped.

Option D is wrong because you still need Route53 to split the traffic.

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

#### NEW QUESTION 3

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

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

**Answer:** C

#### Explanation:

The below diagram shows how this can be done

1) First create a new ELB which will be used to point to the new production changes.

2) Use the Weighted Route policy for Route53 to distribute the traffic to the 2 ELB's based on a 80- 20% traffic scenario. This is the normal case, the % can be changed based on the requirement.

3) Finally when all changes have been tested, Route53 can be set to 100% for the new ELB.

Option A is incorrect because this is a failover scenario and cannot be used for Blue green deployments. In Blue Green deployments, you need to have 2 environments running side by side. Option B is incorrect, because you need to have a production stack with the changes which will run side by side.

Option D is incorrect because this is not a blue green deployment scenario. You cannot control which users will go to the new EC2 instances.

For more information on blue green deployments, please refer to the below document link: from AWS

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

#### NEW QUESTION 4

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 5**

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 CLB 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 6**

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

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

**Answer:** B

**Explanation:**

Elastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot issues. Option A is invalid because this service will monitor all AWS services. Option C and D are invalid since CLB already provides a logging feature. For more information on ELB access logs, please refer to the below document link: from AWS <http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/access-log-collection.html>

**NEW QUESTION 7**

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 fulfils 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 send 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 send 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 8**

During metric analysis, your team has determined that the company's website during peak hours is experiencing response times higher than anticipated. You currently rely on Auto Scaling to make sure that you are scaling your environment during peak windows. How can you improve your Auto Scaling policy to reduce this high response time? Choose 2 answers.

A. Push custom metrics to CloudWatch to monitor your CPU and network bandwidth from your servers, which will allow your Auto Scaling policy to have betterfine-grain insight.

B. IncreaseyourAutoScalinggroup'snumberofmaxservers.

C. Create a script that runs and monitors your servers; when it detects an anomaly in load, it posts to an Amazon SNS topic that triggers Elastic Load Balancing to add more servers to the load balancer.

D. Push custom metrics to CloudWatch for your application that include more detailed information about your web application, such as how many requests it is handling and how many are waiting to be processed.

**Answer: BD**

**Explanation:**

Option B makes sense because maybe the max servers is low hence the application cannot handle the peak load.

Option D helps in ensuring Autoscaling can scale the group on the right metrics.

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

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

**NEW QUESTION 9**

Management has reported an increase in the monthly bill from Amazon Web Services, and they are extremely concerned with this increased cost. Management has asked you to determine the exact cause of this increase. After reviewing the billing report, you notice an increase in the data transfer cost. How can you provide management with a better insight into data transfer use?

A. Update your Amazon CloudWatch metrics to use five-second granularity, which will give better detailed metrics that can be combined with your billing data to pinpoint anomalies.

B. Use Amazon CloudWatch Logs to run a map-reduce on your logs to determine high usage and data transfer.

C. Deliver custom metrics to Amazon CloudWatch per application that breaks down application data transfer into multiple, more specific data points.D- Using Amazon CloudWatch metrics, pull your Elastic Load Balancing outbound data transfer metrics monthly, and include them with your billing report to show which application is causing higher bandwidth usage.

**Answer: C**

**Explanation:**

You can publish your own metrics to CloudWatch using the AWS CLI or an API. You can view statistical graphs of your published metrics with the AWS Management Console.

CloudWatch stores data about a metric as a series of data points. Each data point has an associated time stamp. You can even publish an aggregated set of data points called a statistic set.

If you have custom metrics specific to your application, you can give a breakdown to the management on the exact issue.

Option A won't be sufficient to provide better insights.

Option B is an overhead when you can make the application publish custom metrics Option D is invalid because just the ELB metrics will not give the entire picture

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

<http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.html>

**NEW QUESTION 10**

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

A. Configure the ephemeral policies on your Auto Scaling group to back up on terminate.

B. Configure your Auto Scaling policies to create a snapshot of all ephemeral storage on terminate.

C. Install the CloudWatch Logs Agent on your AMI, and configure CloudWatch Logs Agent to stream your logs.V

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://docsaws.amazon.com/AmazonCloudWatch/latest/logs/WhatsCloudWatchLogs.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 13

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

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

**Answer:** AC

#### Explanation:

Option B is invalid because Cloudtrail is not designed specifically to take in UDP packets  
Option D is invalid because there are already Cloudwatch logs available, so there is no need to have specific logs designed for this.  
You can use Amazon CloudWatch Logs to monitor, store, and access your log files from Amazon Elastic Compute Cloud (Amazon EC2) instances, AWS CloudTrail, and other sources. You can then retrieve the associated log data from CloudWatch Logs. For more information on Cloudwatch logs please refer to the below link: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html> You can use Kinesis to process those logs  
For more information on Amazon Kinesis please refer to the below link: <http://docs.aws.amazon.com/streams/latest/dev/introduction.html>

#### NEW QUESTION 18

You are a DevOps engineer for a company. You have been requested to create a rolling deployment solution that is cost-effective with minimal downtime. How should you achieve this? Choose two answers from the options below

- A. Re-deploy your application using a CloudFormation template to deploy Elastic Beanstalk
- B. Re-deploy with a CloudFormation template, define update policies on Auto Scaling groups in your CloudFormation template
- C. Use UpdatePolicy attribute to specify how CloudFormation handles updates to Auto Scaling Group resource.
- D. After each stack is deployed, tear down the old stack

**Answer:** BC

#### 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. Option A is invalid because it is not efficient to use CloudFormation to use Elastic Beanstalk. Option D is invalid because this is an inefficient process to tear down stacks when there are stack policies available  
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 21

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 provides access to the objects based on the key values generated via the user id. Option A is possible but then becomes a maintenance overhead because of the number of buckets. Option B is invalid because IAM users is not a good security practice. Option D is invalid because SMS tokens are not efficient for this requirement. For more information on the Security Token Service please refer to the below link: <http://docs.aws.amazon.com/STS/latest/APIReference/Welcome.html>

**NEW QUESTION 23**

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

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

**Answer:** A

**Explanation:**

You can use the container\_commands key to execute commands that affect your application source code. Container commands run after the application and web server have been set up and the application version archive has been extracted, but before the application version is deployed. Non-container commands and other customization operations are performed prior to the application source code being extracted. You can use leader\_only to only run the command on a single instance, or configure a test to only run the command when a test command evaluates to true. Leader-only container commands are only executed during environment creation and deployments, while other commands and server customization operations are performed every time an instance is provisioned or updated. Leader-only container commands are not executed due to launch configuration changes, such as a change in the AMI Id or instance type. For more information on customizing containers, please visit the below URL: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/customize-containers-ec2.html>

**NEW QUESTION 24**

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 Elastic Beanstalk-specific JSON file that describes how to deploy a set of Docker containers as an Elastic Beanstalk application. You can use a Dockerrun.aws.json file for a multicontainer Docker environment. Dockerrun.aws.json describes the containers to deploy to each container instance in the environment as well as the data volumes to create on the host instance for the containers to mount. [http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create\\_deploy\\_docker\\_v2config.html](http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_v2config.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 practises please refer to the below link: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html>

**NEW QUESTION 28**

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

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

**Answer: C**

**Explanation:**

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

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

Option A is invalid because this is possible

Option B is invalid because this is used make AWS CloudFormation pause the creation of a stack and wait for a signal before it continues to create the stack

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

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

**NEW QUESTION 30**

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

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

**Answer: B**

**Explanation:**

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

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

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

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

**NEW QUESTION 33**

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

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

**Answer: AB**

**Explanation:**

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

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

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

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

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

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

**NEW QUESTION 37**

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 40**

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 45**

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

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

**Answer:** D

**Explanation:**

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

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

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

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

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

**NEW QUESTION 49**

You need to deploy a Node.js application and do not have any experience with AWS. Which deployment method will be the simplest for you to deploy?

- A. AWS Elastic Beanstalk
- B. AWS CloudFormation
- C. AWS EC2
- D. AWS OpsWorks

**Answer:** A

**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.

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

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

**NEW QUESTION 53**

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

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



**Answer:** D

**Explanation:**

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

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

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

**NEW QUESTION 56**

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 60**

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

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

**Answer:** A

**Explanation:**

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

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

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

**NEW QUESTION 61**

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, it's 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 AMI 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/AWSCC2/latest/UserGuide/AMIs.html>

#### NEW QUESTION 66

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

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

**Answer:** A

#### Explanation:

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

This is ideal for Large database workloads, such as:

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

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

- <http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/CBSVolumeTypes.html>

#### NEW QUESTION 68

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. GameIDastherange/onlykey.

**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:

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 73

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

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

**Answer:** C

#### Explanation:

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

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

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

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

- <http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/CBSEncryption.html>

#### NEW QUESTION 76

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

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

**Answer:** B

**Explanation:**

The CBS Snapshots are stored in S3, so if you have an scripts which deploy CC2 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/AWSSCC2/latest/UserGuide/CBSSnapshots.htm>

**NEW QUESTION 79**

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

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

**Answer: C**

**Explanation:**

A tag is a label that you or AWS assigns to an AWS resource. Each tag consists of a key and a value. A key can have more than one value. You can use tags to organize your resources, and cost allocation tags to track your AWS costs on a detailed level. After you activate cost allocation tags, AWS uses the cost allocation tags to organize your resource costs on your cost allocation report, to make it easier for you to categorize and track your AWS costs. AWS provides two types of cost allocation tags, an AWS-generated tag and user-defined tags. AWS defines, creates, and applies the AWS-generated tag for you, and you define, create, and apply user-defined tags. You must activate both types of tags separately before they can appear in Cost Explorer or on a cost allocation report.

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

**NEW QUESTION 80**

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

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

**Answer: D**

**Explanation:**

A placement group is a logical grouping of instances within a single Availability Zone. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. To provide the lowest latency, and the highest packet-per-second network performance for your placement group, choose an instance type that supports enhanced networking. For more information on Placement Groups, please visit the below URL: <http://docs.aws.amazon.com/AWSSCC2/latest/UserGuide/placement-groups.html>

**NEW QUESTION 84**

You need to deploy a new application version to production. Because the deployment is high-risk, you need to roll the new version out to users over a number of hours, to make sure everything is working correctly. You need to be able to control the proportion of users seeing the new version of the application down to the percentage point. You use ELB and EC2 with Auto Scaling Groups and custom AMIs with your code pre-installed assigned to Launch Configurations. There are no database-level changes during your deployment. You have been told you cannot spend too much money, so you must not increase the number of EC2 instances much at all during the deployment, but you also need to be able to switch back to the original version of code quickly if something goes wrong. What is the best way to meet these requirements?

- A. Create a second ELB, Auto Scaling Launch Configuration, and Auto Scaling Group using the Launch Configuration.
- B. Create AMIs with all code pre-installed.
- C. Assign the new AMI to the second Auto Scaling Launch Configuration.
- D. Use Route53 Weighted Round Robin Records to adjust the proportion of traffic hitting the two ELBs.
- E. Use the Blue-Green deployment method to enable the fastest possible rollback if needed.
- F. Create a full second stack of instances and cut the DNS over to the new stack of instances, and change the DNS back if a rollback is needed.
- G. Create AMIs with all code pre-installed.
- H. Assign the new AMI to the Auto Scaling Launch Configuration, to replace the old one.
- I. Gradually terminate instances running the old code (launched with the old Launch Configuration) and allow the new AMIs to boot to adjust the traffic balance to the new code.
- J. On rollback, reverse the process by doing the same thing, but changing the AMI on the Launch Config back to the original code.
- K. Migrate to use AWS Elastic Beanstalk.
- L. Use the established and well-tested Rolling Deployment setting AWS provides on the new Application Environment, publishing a zip bundle of the new code and adjusting the wait period to spread the deployment over time.
- M. Re-deploy the old code bundle to rollback if needed.

**Answer: A**

**Explanation:**

This is an example of a Blue Green Deployment.

You can shift traffic all at once or you can do a weighted distribution. 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

For more information on Blue Green Deployments, please visit the below URL:

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

#### NEW QUESTION 87

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

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

**Answer:** A

#### Explanation:

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

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

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

#### NEW QUESTION 91

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. AWSCloudFormation
- B. AWSOpsWorks
- 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 92

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

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

**Answer:** B

#### Explanation:

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



#### Connector

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

- Your end users and IT administrators can use their existing corporate credentials to log on to AWS applications such as Amazon Workspaces, Amazon WorkDocs, or Amazon WorkMail.
- You can manage AWS resources like Amazon EC2 instances or Amazon S3 buckets through 1AM role-based access to the AWS Management Console.
- You can consistently enforce existing security policies (such as password expiration, password history, and account lockouts) whether users or IT administrators are accessing resources in your on-premises infrastructure or in the AWS Cloud.
- You can use AD Connector to enable multi-factor authentication by integrating with your existing RADIUS-based MFA infrastructure to provide an additional layer of security when users access AWS applications.

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

- [http://docs.aws.amazon.com/directoryservice/latest/admin-guide/directory\\_ad\\_connector.htm](http://docs.aws.amazon.com/directoryservice/latest/admin-guide/directory_ad_connector.htm)

#### NEW QUESTION 96

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

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

**Answer: D**

#### Explanation:

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

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

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

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

#### NEW QUESTION 100

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

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

**Answer: A**

#### Explanation:

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

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

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

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

#### NEW QUESTION 102

Your serverless architecture using AWS API Gateway, AWS Lambda, and AWS DynamoDB experienced a large increase in traffic to a sustained 3000 requests per second, and dramatically increased in failure rates. Your requests, during normal operation, last 500 milliseconds on average. Your DynamoDB table did not exceed 50% of provisioned throughput, and Table primary keys are designed correctly. What is the most likely issue?

- A. Your API Gateway deployment is throttling your requests.
- B. Your AWS API Gateway Deployment is bottlenecking on request (deserialization).
- C. You did not request a limit increase on concurrent Lambda function executions.
- D. You used Consistent Read requests on DynamoDB and are experiencing semaphore lock.

**Answer: C**

#### Explanation:

Every Lambda function is allocated with a fixed amount of specific resources regardless of the memory allocation, and each function is allocated with a fixed amount of code storage per function and per account.

By default, AWS Lambda limits the total concurrent executions across all functions within a given region to 1000.

For more information on Concurrent executions, please visit the below URL: <http://docs.aws.amazon.com/lambda/latest/dg/concurrent-executions.html>

#### NEW QUESTION 104

You currently have an application deployed via Elastic Beanstalk. You are now deploying a new application and have ensured that Elastic beanstalk has detached the current instances and deployed and reattached new instances. But the new instances are still not receiving any sort of traffic. Why is this the case?

- A. The instances are of the wrong AMI, hence they are not being detected by the ELB.
- B. It takes time for the ELB to register the instances, hence there is a small timeframe before your instances can start receiving traffic
- C. You need to create a new Elastic Beanstalk application, because you cannot detach and then reattach instances to an ELB within an Elastic Beanstalk application
- D. The instances needed to be reattached before the new application version was deployed

**Answer: B**

**Explanation:**

Before the EC2 Instances can start receiving traffic, they will be checked via the health checks of the CLB. Once the health checks are successful, the EC2 Instance

will change its state to InService and then the EC2 Instances can start receiving traffic. 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 108**

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

- A. All at Once
- B. Rolling Deployments
- C. Rolling with additional batch
- D. Immutable

**Answer: B**

**Explanation:**

The AWS documentation mentions

AWS Elastic Beanstalk provides several options for how deployments are processed, including deployment policies (All at once, Rolling, Rolling with additional batch,

and Immutable) and options that let you configure batch size and health check behavior during deployments. By default, your environment uses rolling deployments

if you created it with the console or EB CLI, or all at once deployments if you created it with a different client (API, SDK or AWS CLI).

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

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

**NEW QUESTION 109**

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 110**

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

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

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

**Answer: C**

**Explanation:**

The below diagram from the AWS documentation shows the deployment steps

For more information on the deployment steps please refer to the below link:

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

**NEW QUESTION 113**

You have an Opsworks 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 116**

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

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

**Answer:** A

**Explanation:**

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

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

**NEW QUESTION 121**

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

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

**Answer:** C

**Explanation:**

The AWS Documentation specifies

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

You can use a Dockerrun.aws.json file for a multicontainer Docker environment.

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

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

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

**NEW QUESTION 126**

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

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

**Answer:** C

**Explanation:**

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

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

**NEW QUESTION 127**

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

Desired capacity - 2 Maximum capacity - 2

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

Which of the below solutions will meet this demand.

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

**Answer:** C

**Explanation:**

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

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

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

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

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

**NEW QUESTION 128**

Your company has the requirement to set up instances running as part of an Autoscaling Group. Part of the requirement is to use Lifecycle hooks to setup custom based software's and do the necessary configuration on the instances. The time required for this setup might take an hour, or might finish before the hour is up. How should you setup lifecycle hooks for the Autoscaling Group. Choose 2 ideal actions you would include as part of the lifecycle hook.

- A. Configure the lifecycle hook to record heartbeat
- B. If the hour is up, restart the timeout period.
- C. Configure the lifecycle hook to record heartbeat
- D. If the hour is up, choose to terminate the current instance and start a new one
- E. If the software installation and configuration is complete, then restart the time period.
- F. If the software installation and configuration is complete, then send a signal to complete the launch of the instance.

**Answer:** AD

**Explanation:**

The AWS Documentation provides the following information on lifecycle hooks

By default, the instance remains in a wait state for one hour, and then Auto Scaling continues the launch or terminate process (Pending: Proceed or Terminating: Proceed). If you need more time, you can restart the timeout period by recording a heartbeat. If you finish before the timeout period ends, you can complete the lifecycle action, which continues the launch or termination process

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

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

**NEW QUESTION 129**

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

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

**Answer:** AD

**Explanation:**

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

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

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

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

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

**NEW QUESTION 133**

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

- A. Create a Cloudformation template with the same resources as those in the Elastic beanstalk environmen
- B. If the deployment fails, deploy the Cloudformation template.



- C. Use Rolling updates in Elastic Beanstalk so that if the deployment fails, the rolling updates feature would roll back to the last deployment.
- D. Create another parallel environment in elastic beanstal
- E. Use the Swap URL feature.
- F. Create another parallel environment in elastic beanstal
- G. Create a new Route53 Domain name for the new environment and release that url to the users.

**Answer:** C

**Explanation:**

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

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

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

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

**NEW QUESTION 135**

You have a legacy application running that uses an m4.large instance size and cannot scale with Auto Scaling, but only has peak performance 5% of the time. This is a huge waste of resources and money so your Senior Technical Manager has set you the task of trying to reduce costs while still keeping the legacy application running as it should. Which of the following would best accomplish the task your manager has set you? Choose the correct answer from the options below

- A. Use a T2burstable performance instance.
- B. Use a C4.large instance with enhanced networking.
- C. Use two t2.nano instances that have single Root I/O Visualization.
- D. Use t2.nano instance and add spot instances when they are required.

**Answer:** A

**Explanation:**

The aws documentation clearly indicates using T2 CC2 instance types for those instances which don't use CPU that often.

T2

T2 instances are Burstable Performance Instances that provide a baseline level of CPU performance with the ability to burst above the baseline.

T2 Unlimited instances can sustain high CPU performance for as long as a workload needs it. For most general-purpose workloads, T2 Unlimited instances will provide ample performance without any additional charges. If the instance needs to run at higher CPU utilization for a prolonged period, it can also do so at a flat additional charge of 5 cents per vCPU-hour.

The baseline performance and ability to burst are governed by CPU Credits. T2 instances receive CPU Credits continuously at a set rate depending on the instance size, accumulating CPU Credits when they are idle, and consuming CPU credits when they are active. T2 instances are a good choice for a variety of general-purpose workloads including micro-services, low-latency interactive applications, small and medium databases, virtual desktops, development, build and stage environments, code repositories, and product prototypes. For more information see Burstable Performance Instances.

For more information on F\_C2 instance types please see the below link: <https://aws.amazon.com/ec2/instance-types/>

**NEW QUESTION 138**

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

- A. During the install and game configuration process, have each user create an 1AM credential and assign the 1AM user to a group with proper permissions to communicate with DynamoDB.
- B. Create an 1AM group that only gives access to your application and to the DynamoDB table
- C. Then, when writing to DynamoDB, simply include the unique device ID to associate the data with that specific user.
- D. Create an 1AM role with the proper permission policy to communicate with the DynamoDB tabl
- E. Use web identity federation, which assumes the 1AM role using AssumeRoleWithWebIdentity, when the user signs in, granting temporary security credentials using STS.
- F. Create an Active Directory server and an AD user for each mobile application use
- G. When the user signs in to the AD sign-on, allow the AD server to federate using SAML 2.0 to 1AM and assign a role to the AD user which is the assumed with AssumeRoleWithSAML

**Answer:** C

**Explanation:**

Answer - C

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

For more information on 1AM policies please refer to the below link:

[http://docs.aws.amazon.com/IAM/latest/UserGuide/access\\_policies.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html)

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

The AWS documentation mentions the following

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

For more information on web identity federation please refer to the below link: [http://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles\\_providers\\_oidc.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html)

**NEW QUESTION 139**

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

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

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

**Answer:** A

**Explanation:**

The AWS Documentation mentions

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

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

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

**NEW QUESTION 144**

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

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

**Answer:** BC

**Explanation:**

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

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

The AWS Documentation mentions

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

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

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

**NEW QUESTION 146**

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

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

**Answer:** B

**Explanation:**

The AWS documentation mentions

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

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

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

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

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

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

**NEW QUESTION 149**

A custom script needs to be passed to a new Amazon Linux instances created in your Auto Scalinggroup. Which feature allows you to accomplish this?

- A. User data  
B. EC2Config service  
C. IAM roles  
D. AWSConfig

**Answer:** A

**Explanation:**

When you configure an instance during creation, you can add custom scripts to the User data section. So in Step 3 of creating an instance, in the Advanced Details section, we can enter custom scripts in the User Data section. The below script installs Perl during the instance creation of the EC2 instance.

For more information on user data please refer to the URL:

- <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-metadata.html>

**NEW QUESTION 154**

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

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

**Answer:** D

**Explanation:**

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

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

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

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

**NEW QUESTION 155**

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 157**

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

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

**Answer:** B

**Explanation:**

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

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

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

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

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

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

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

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

**NEW QUESTION 159**

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

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

**Answer:** ABC

**NEW QUESTION 161**

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. Haveone single dashboard to report metrics to CloudWatch from different region
- D. Thisis not possible

**Answer: C**

**Explanation:**

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

CC2 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 165**

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

- A. UsingS3 Server Side Encryption (SSE) to store the information
- B. Enable SSLtermination on the ELB C EnablingProxy ProtocolD- Enablingsticky sessions on your load balancer

**Answer: B**

**Explanation:**

Please note the keyword "NOT" in the question.

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

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

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

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

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

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

Requirements

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

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

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

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

<http://docsaws.amazon.com/elasticloadbalancing/latest/classic/elb-https-load-balancers.html>

<https://aws.amazon.com/blogs/aws/elastic-load-balancer-support-for-ssl-termination/>

**NEW QUESTION 168**

You are currently planning on using Autoscaling to launch instances which have an application installed. Which of the following methods will help ensure the instances are up and running in the shortest span of time to take in traffic from the users?

- A. Loginto each instance and install the software.
- B. UseUserData to launch scripts to install the software.
- C. UseDocker containers to launch the software.
- D. UseAMI's which already have the software installed.

**Answer: D**

**Explanation:**

The AM I will be the fatest because it will already have the software installed. You can customize the instance that you launch from a public AMI and then save that configuration as a custom AMI for your own use. Instances that you launch from your AMI use all the custom izations that you've made.

For more information on AMI'S please refer to the below link <http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/AMIs.html>

**NEW QUESTION 169**

Which of the following are advantages of using AWS CodeCommit over hosting your own source code repository system?

- A. Reduction in hardware maintenance costs
- B. Reduction in fees paid over licensing
- C. No specific restriction on files andbranches
- D. All of the above

**Answer: D**

**Explanation:**

The AWS Documentation mentions the following on CodeCommit

Self-hosted version control systems have many potential drawbacks, including: Expensive per-developer licensing fees.

High hardware maintenance costs. High support staffing costs.

Limits on the amount and types of files that can be stored and managed.

Limits on the number of branches, the amount of version history, and other related metadata that can be stored. For more information on CodeCommit please refer to the below link



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

### NEW QUESTION 173

Which of the following is not a supported platform on Elastic Beanstalk?

- A. PackerBuilder
- B. Go
- C. Nodejs
- D. JavaSE
- E. Kubernetes

**Answer:** E

#### Explanation:

Answer-C

Below is the list of supported platforms

- \*Packer Builder
- \*Single Container Docker
- \*Multicontainer Docker
- \*Preconfigured Docker
- \*Go
- \*Java SE
- \*Java with Tomcat
- \*NET on Windows Server with IIS
- \*Nodejs
- \*PHP
- \*Python
- \*Ruby

For more information on the supported platforms please refer to the below link

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/concepts.platforms.html>

### NEW QUESTION 176

You have a number of Cloudformation stacks in your IT organization. Which of the following commands will help see all the cloudformation stacks which have a completed status?

- A. describe-stacks
- B. list-stacks
- C. stacks-complete
- D. list-templates

**Answer:** B

#### Explanation:

The following is the description of the list-stacks command

Returns the summary information for stacks whose status matches the specified StackStatusFilter.

Summary information for stacks that have been deleted is kept for 90 days after the stack is deleted. If no stack-status-filter is specified, summary information for all stacks is returned (including existing stacks and stacks that have been deleted).

For more information on the list-stacks command please visit the below link <http://docs.aws.amazon.com/cli/latest/reference/cloudformation/list-stacks.html>

### NEW QUESTION 180

When one creates an encrypted EBS volume and attach it to a supported instance type, which of the following data types are encrypted?

Choose 3 answers from the options below

- A. Data at rest inside the volume
- B. All data copied from the EBS volume to S3
- C. All data moving between the volume and the instance
- D. All snapshots created from the volume

**Answer:** ACD

#### Explanation:

This is clearly given in the AWS documentation. Amazon EBS Encryption

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

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

For more information on EBS encryption, please refer to the below URL <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html>

### NEW QUESTION 185

You are in charge of designing a number of Cloudformation templates for your organization. You need to ensure that no one can accidentally update the production-based resources on the stack during a stack update. How can this be achieved in the most efficient way?

- A. Create tags for the resources and then create IAM policies to protect the resources.
- B. Use a Stack-based policy to protect the production-based resources.
- C. Use S3 bucket policies to protect the resources.
- D. Use MFA to protect the resources

**Answer:** B

**Explanation:**

The AWS Documentation mentions

When you create a stack, all update actions are allowed on all resources. By default, anyone with stack update permissions can update all of the resources in the stack. During an update, some resources might require an interruption or be completely replaced, resulting in new physical IDs or completely new storage. You can prevent stack resources from being unintentionally updated or deleted during a stack update by using a stack policy. A stack policy is a JSON document that defines the update action1.-; that can be performed on designated resources.

For more information on protecting stack resources, please visit the below url <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/protect-stack-resources.html>

**NEW QUESTION 189**

You are in charge of designing a number of CloudFormation templates for your organization. You are required to make changes to stack resources every now and then based on the requirement. How can you check the impact of the change to resources in a cloudformation stack before deploying changes to the stack?

- A. There is no way to control this
- B. You need to check for the impact beforehand.
- C. Use CloudFormation change sets to check for the impact to the changes.
- D. Use CloudFormation Stack Policies to check for the impact to the changes.
- E. Use CloudFormation Rolling Updates to check for the impact to the changes.

**Answer:** B

**Explanation:**

The AWS Documentation mentions

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

CloudFormation console, AWS CLI, or AWS CloudFormation API.

For more information on CloudFormation change sets, please visit the below url <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks-changesets.html>

**NEW QUESTION 193**

Which of the following is false when it comes to using the Elastic Load balancer with Opsworks stacks?

- A. You can attach only one load balancer to a layer.
- B. A Classic Load Balancer can span across AWS OpsWorks Stacks layers.
- C. Each load balancer can handle only one layer.
- D. You need to create the load balancer before hand and then attach it to the Opsworkstack.

**Answer:** B

**Explanation:**

The AWS Documentation mentions the following

To use Elastic Load Balancing with a stack, you must first create one or more load balancers in the same region by using the Elastic Load Balancing console, CLI, or API. You should be aware of the following:

You can attach only one load balancer to a layer. Each load balancer can handle only one layer.

AWS OpsWorks Stacks does not support Application Load Balancer. You can only use Classic Load Balancer with AWS OpsWorks Stacks. For more information on Elastic Load Balancer with Opswork, please visit the below url <http://docs.aws.amazon.com/opsworks/latest/userguide/layers-elb.html>

**NEW QUESTION 198**

You currently have a set of instances running on your Opswork stacks. You need to install security updates on these servers. What does AWS recommend in terms of how the security updates should be deployed?

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. Create a new Opswork stack with the new instances.
- D. On Linux-based instances in Chef 11.10 or older stacks, run the UpdateDependencies stack command.
- E. Create a cloudformation template which can be used to replace the instances.

**Answer:** AC

**Explanation:**

The AWS Documentation mentions the following

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.

For more information on Opswork updates, please visit the below url • <http://docs.aws.amazon.com/opsworks/latest/userguide/best-practices-updates.html>

**NEW QUESTION 203**

You are creating a Cloudformation template in which UserData is going to be passed to underlying EC2 Instance. Which of the below functions is normally used to pass data to the UserData section in the Cloudformation template?

- A. "UserData": { "Fn::Base64":{
- B. "UserData": < "Fn::Ref":{
- C. "UserData": { "Fn::GetAtt":{
- D. "UserData": { "Fn::FindInMap":{

**Answer:** A

**Explanation:**

The AWS Documentation mentions

The intrinsic function Fn::Base64 returns the Base64 representation of the input string. This function is typically used to pass encoded data to Amazon EC2 instances by way of the User Data property.

For more information on the Fn::Base64 function, please visit the below URL: [http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-base64.htm](http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-base64.html) l

**NEW QUESTION 204**

Which of the following are the basic stages of a CI/CD Pipeline. Choose 3 answers from the options below

- A. SourceControl
- B. Build
- C. Run
- D. Production

**Answer:** ABD

**Explanation:**

The below diagram shows the stages of a typical CI/CD pipeline

For more information on AWS Continuous Integration, please visit the below URL: <https://d.awsstatic.com/whitepapers/DevOps/practicing-continuous-integration-continuous-delivery-on-AWS.pdf>

**NEW QUESTION 209**

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 be 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 210**

You have just developed a new mobile application that handles analytics workloads on large scale datasets that are stored on Amazon Redshift. Consequently, the application needs to access Amazon Redshift tables. Which of the below methods would be the best, both practically and security-wise, to access the tables?

Choose the correct answer from the options below

- A. Create an IAM user and generate encryption keys for that use
- B. Create a policy for RedShift read-only access
- C. Embed the keys in the application.
- D. Create an HSM client certificate in Redshift and authenticate using this certificate.
- E. Create a RedShift read-only access policy in IAM and embed those credentials in the application.
- F. User roles that allow a web identity federated user to assume a role that allows access to the RedShift table by providing temporary credentials.

**Answer:** D

**Explanation:**

For access to any AWS service, the ideal approach for any application is to use Roles. This is the first preference. Hence option A and C are wrong.

For more information on IAM policies please refer to the below link: [http://docs.aws.amazon.com/IAM/latest/UserGuide/access\\_policies.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html)

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

"When you write such an app, you'll make requests to AWS services that must be signed with an AWS access key. However, we strongly recommend that you do not embed or distribute long-term AWS credentials with apps that a user downloads to a device, even in an encrypted store. Instead, build your app so that it requests temporary AWS security credentials dynamically when needed using web identity federation. The supplied temporary credentials map to an AWS role that has only

the permissions needed to perform the tasks required by the mobile app".

For more information on web identity federation please refer to the below link: [http://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles\\_providers\\_oidc.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html)

**NEW QUESTION 214**

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

**Answer:** B

**Explanation:**

Since the AMI will get all the data from S3 as snapshots, always ensure the volume prewarmed before it is set for the load test.

For more information on benchmarking procedures please see the below link:

- [http://docs^ws.amazon.com/AWSSCC2/latest/UserGuide/berK;hmark\\_prooedures.html](http://docs.aws.amazon.com/AWSSCC2/latest/UserGuide/berK;hmark_prooedures.html)

**NEW QUESTION 215**

Your social media marketing application has a component written in Ruby running on AWS Elastic Beanstalk. This application component posts messages to social media sites in support of various marketing campaigns. Your management now requires you to record replies to these social media messages to analyze the effectiveness of the marketing campaign in comparison to past and future efforts. You've already developed a new application component to interface with the social media site APIs in order to read the replies. Which process should you use to record the social media replies in a durable data store that can be accessed at any time for analytics of historical data?

- A. Deploythe new application component in an Auto Scaling group of Amazon EC2 instances,read the data from the social media sites, store it with Amazon Elastic BlockStore, and use AWS Data Pipeline to publish it to Amazon Kinesis for analytics.
- B. Deploythe new application component as an Elastic Beanstalk application, read thedata from the social media sites, store it in DynamoDB, and use Apache Hivewith Amazon Elastic MapReduce for analytics.
- C. Deploythe new application component in an Auto Scaling group of Amazon EC2 instances,read the data from the social media sites, store it in Amazon Glacier, and useAWS Data Pipeline to publish it to Amazon RedShift for analytics.
- D. Deploythe new application component as an Amazon Elastic Beanstalk application, readthe data from the social media site, store it with Amazon Elastic Block store,and use Amazon Kinesis to stream the data to Amazon Cloud Watch for analytics

**Answer:** B

**Explanation:**

The AWS Documentation mentions the below

Amazon DynamoDB is a fast and flexible NoSQL database sen/ice for all applications that need consistent, single-digit millisecond latency at any scale. It is a fully managed cloud database and supports both document and key-value store models. Its flexible data model, reliable performance, and automatic scaling of throughput capacity, makes it a great fit for mobile, web, gaming, ad tech, IoT, and many other applications.

For more information on AWS DynamoDB please see the below link:

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

**NEW QUESTION 218**

Which of the following can be configured as targets for Cloudwatch Events. Choose 3 answers from the options given below

- A. AmazonEC2 Instances
- B. AWSLambda Functions
- C. AmazonCodeCommit
- D. AmazonECS Tasks

**Answer:** ABD

**Explanation:**

The AWS Documentation mentions the below

You can configure the following AWS sen/ices as targets for Cloud Watch Events

For more information on Cloudwatch events please see the below link:

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

**NEW QUESTION 222**

You are a Devops Engineer for your company. You are in charge of an application that uses EC2, ELB and Autoscaling. You have been requested to get the ELB access logs. When you try to access the logs, you can see that nothing has been recorded in S3. Why is this the case?

- A. Youdon't have the necessary access to the logs generated by ELB.
- B. Bydefault ELB access logs are disabled.
- C. TheAutoscaling service is not sending the required logs to ELB
- D. TheEC2 Instances are not sending the required logs to ELB

**Answer:** B

**Explanation:**

The AWS Documentation mentions

Access logging is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logging for your load balancer. Clastic 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 L~LB access logs please see the below link:

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

**NEW QUESTION 223**

Which of the following is a container for metrics in Cloudwatch?



- A. MetricCollection
- B. Namespaces
- C. Packages
- D. Locale

**Answer:** B

**Explanation:**

The AWS Documentation mentions the following

Cloud Watch namespaces are containers for metrics. Metrics in different namespaces are isolated from each other, so that metrics from different applications are not mistakenly aggregated into the same statistics. All AWS services that provide Amazon Cloud Watch data use a namespace string, beginning with "AWS/".

When

you create custom metrics, you must also specify a namespace as a container for custom metrics. The following services push metric data points to Cloud Watch.

For more information on Cloudwatch namespaces, please visit the below URL: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/aws-namespaces.html>

**NEW QUESTION 227**

You work as a Devops Engineer for your company. There are currently a number of environments hosted via Elastic beanstalk. There is a requirement to ensure to ensure that the rollback time for a new version application deployment is kept to a minimal. Which elastic beanstalk deployment method would fulfil this requirement ?

- A. Rollingwith additional batch
- B. AllatOnce
- C. Blue/Green
- D. Rolling

**Answer:** C

**Explanation:**

The below table from the AWS documentation shows that the least amount of time is spent in rollbacks when it comes to Blue Green deployments. This is because the only thing that needs to be done is for URL's to be swapped.

For more information on Elastic beanstalk deployment strategies, please visit the below URL: [http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.deploy-existing-version, html](http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.deploy-existing-version.html)

**NEW QUESTION 228**

.....

## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### DOP-C01 Practice Exam Features:

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

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