

## DOP-C01 Dumps

### AWS Certified DevOps Engineer- Professional

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



**NEW QUESTION 1**

You have an application which consists of EC2 instances in an Auto Scaling group. Between a particular time frame every day, there is an increase in traffic to your website. Hence users are complaining of a poor response time on the application. You have configured your Auto Scaling group to deploy one new EC2 instance when CPU utilization is greater than 60% for 2 consecutive periods of 5 minutes. What is the least cost-effective way to resolve this problem?

- A. Decrease the consecutive number of collection periods
- B. Increase the minimum number of instances in the Auto Scaling group
- C. Decrease the collection period to ten minutes
- D. Decrease the threshold CPU utilization percentage at which to deploy a new instance

**Answer: B**

**Explanation:**

If you increase the minimum number of instances, then they will be running even though the load is not high on the website. Hence you are incurring cost even though there is no need.

All of the remaining options are possible options which can be used to increase the number of instances on a high load.

For more information on On-demand scaling, please refer to the below link: <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scale-based-on-demand.html>

Note: The tricky part where the question is asking for 'least cost effective way'. You got the design consideration correctly but need to be careful on how the question is phrased.

**NEW QUESTION 2**

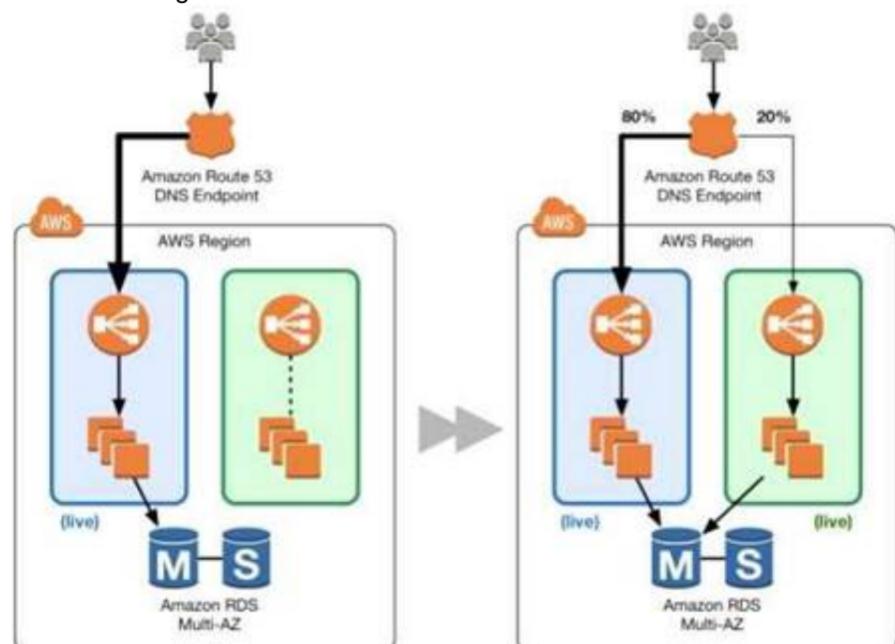
Your application is currently running on Amazon EC2 instances behind a load balancer. Your management has decided to use a Blue/Green deployment strategy. How should you implement this for each deployment?

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

**Answer: C**

**Explanation:**

The below diagram shows how this can be done



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

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

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

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

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

**NEW QUESTION 3**

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

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

**Answer: D**

**Explanation:**

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

**NEW QUESTION 4**

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

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

**Answer:** D

**Explanation:**

Elastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot issues. Access logging is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logging for your load balancer, Elastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify. You can disable access logging at any time. For more information on 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 5**

Which Auto Scaling process would be helpful when testing new instances before sending traffic to them, while still keeping them in your Auto Scaling Group?

- A. Suspend the process AZ Rebalance
- B. Suspend the process Health Check
- C. Suspend the process Replace Unhealthy
- D. Suspend the process AddToLoadBalancer

**Answer:** D

**Explanation:**

If you suspend AddToLoadBalancer, Auto Scaling launches the instances but does not add them to the load balancer or target group. If you resume the AddToLoadBalancer process, Auto Scaling resumes adding instances to the load balancer or target group when they are launched. However, Auto Scaling does not add the instances that were launched while this process was suspended. You must register those instances manually. Option A is invalid because this just balances the number of EC2 instances in the group across the Availability Zones in the region. Option B is invalid because this just checks the health of the instances. Auto Scaling marks an instance as unhealthy if Amazon EC2 or Elastic Load Balancing tells Auto Scaling that the instance is unhealthy. Option C is invalid because this process just terminates instances that are marked as unhealthy and later creates new instances to replace them. For more information on process suspension, please refer to the below document link: from AWS <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resume-processes.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**

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

- A. Using AWS CloudFormation, create a CloudWatch Logs LogGroup and send the operating system and application logs of interest using the CloudWatch Logs

Agent.

B. Using AWS CloudFormation and configuration management, set up remote logging to send events via UDP packets to CloudTrail.

C. Using configuration management, set up remote logging to send events to Amazon Kinesis and insert these into Amazon CloudSearch or Amazon Redshift, depending on available analytic tools.

D. Using AWS CloudFormation, merge the application logs with the operating system logs, and use IAM Roles to allow both teams to have access to view console output from Amazon EC2.

**Answer:** AC

**Explanation:**

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

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

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

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

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

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

**NEW QUESTION 8**

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

A. Create an Amazon S3 bucket per user, and use your application to generate the S3 URI for the appropriate content.

B. Use AWS Identity and Access Management (IAM) user accounts as your application-level user database, and offload the burden of authentication from your application code.

C. Authenticate your users at the application level, and use AWS Security Token Service (STS) to grant token-based authorization to S3 objects.

D. Authenticate your users at the application level, and send an SMS token message to the user

E. Create an Amazon S3 bucket with the same name as the SMS message token, and move the user's objects to that bucket.

F. Use a key-based naming scheme comprised from the user IDs for all user objects in a single Amazon S3 bucket.

**Answer:** CE

**Explanation:**

The AWS Security Token Service (STS) is a web service that enables you to request temporary, limited-privilege credentials for AWS Identity and Access Management (IAM) users or for users that you authenticate (federated users). The token can then be used to grant access to the objects in S3.

You can then provide access to the objects based on the key values generated via the user id. Option A is possible but then becomes a maintenance overhead because of the number of buckets. Option B is invalid because IAM users is not a good security practice.

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

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

**NEW QUESTION 9**

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

A. Add Amazon RDS DB read replicas, and have your application direct read queries to them.

B. Add your Amazon RDS DB instance to an Auto Scaling group and configure your CloudWatch metric-based on CPU utilization.

C. Use an Amazon SQS queue to throttle data going to the Amazon RDS DB instance.

D. Use ElastiCache in front of your Amazon RDS DB to cache common queries.

E. Shard your data set among multiple Amazon RDS DB instances.

F. Enable Multi-AZ for your Amazon RDS DB instance.

**Answer:** ADE

**Explanation:**

Amazon RDS Read Replicas provide enhanced performance and durability for database (DB) instances. This replication feature makes it easy to elastically scale out

beyond the capacity constraints of a single DB Instance for read-heavy database workloads. You can create one or more replicas of a given source DB Instance and

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

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

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

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

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

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

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

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

Option B is not an ideal way to scale a database

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

**NEW QUESTION 10**

You have a current CloudFormation template defined in AWS. You need to change the current alarm threshold defined in the CloudWatch alarm. How can you

achieve this?

- A. Currently, there is no option to change what is already defined in Cloudformation templates.
- B. Update the template and then update the stack with the new templat
- C. Automatically all resources will be changed in the stack.
- D. Update the template and then update the stack with the new templat
- E. Only those resources that need to be changed will be change
- F. All other resources which do not need to be changed will remain as they are.
- G. Delete the current cloudformation templat
- H. Create a new one which will update the current resources.

**Answer: C**

**Explanation:**

Option A is incorrect because Cloudformation templates have the option to update resources.

Option B is incorrect because only those resources that need to be changed as part of the stack update are actually updated.

Option D is incorrect because deleting the stack is not the ideal option when you already have a change option available.

When you need to make changes to a stack's settings or change its resources, you update the stack instead of deleting it and creating a new stack. For example, if you

have a stack with an EC2 instance, you can update the stack to change the instance's AMI ID.

When you update a stack, you submit changes, such as new input parameter values or an updated template. AWS CloudFormation compares the changes you submit with the current state of your stack and updates only the changed resources

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

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

**NEW QUESTION 10**

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

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

**Answer: B**

**Explanation:**

During the AMI-creation process, Amazon EC2 creates snapshots of your instance's root volume and any other EBS volumes attached to your instance

New EBS volumes receive their maximum performance the moment that they are available and do not require initialization (formerly known as pre-warming).

However, storage blocks on volumes that were restored from snapshots must be initialized (pulled

down from Amazon S3 and written to the volume) before you can access the block. This preliminary action takes time and can cause a significant increase in the latency of an I/O operation the first time each block is accessed. For most applications, amortizing this cost over the lifetime of the volume is acceptable.

Option A is invalid because block sizes are predetermined and should not be randomly selected. Option C is invalid because this is part of continuous integration and hence volumes can be destroyed after the test and hence there should not be snapshots created unnecessarily

Option D is invalid because the encryption is a security feature and not part of load tests normally. For more information on EBS initialization please refer to the below link:

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

**NEW QUESTION 15**

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

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

**Answer: BDE**

**Explanation:**

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

CloudWatch Logs uses your log data for monitoring; so, no code changes are required. For example, you can monitor application logs for specific literal terms (such as "NullPointerException") or count the number of occurrences of a literal term at a particular position in log data (such as "404" status codes in an Apache access log). When the term you are searching for is found, CloudWatch Logs reports the data to a CloudWatch metric that you specify. Log data is encrypted while in transit and while it is at rest

For more information on Cloudwatch logs please refer to the below link: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html>

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

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

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

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

**NEW QUESTION 19**

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 Creation Policy 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 Creation Policy 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 23**

Your development team wants account-level access to production instances in order to do live debugging of a highly secure environment. Which of the following should you do?

- A. Place the credentials provided by Amazon Elastic Compute Cloud (EC2) into a secure Amazon Simple Storage Service (S3) bucket with encryption enable
- B. Assign AWS Identity and Access Management (IAM) users to each developer so they can download the credentials file.
- C. Place an internally created private key into a secure S3 bucket with server-side encryption using customer keys and configuration management, create a service account on all the instances using this private key, and assign IAM users to each developer so they can download the file.
- D. Place each developer's own public key into a private S3 bucket, use instance profiles and configuration management to create a user account for each developer on all instances, and place the user's public keys into the appropriate account
- E. ^/
- F. Place the credentials provided by Amazon EC2 onto an MFA encrypted USB drive, and physically share it with each developer so that the private key never leaves the office.

**Answer: C**

**Explanation:**

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

A private S3 bucket can be created for each developer, the keys can be stored in the bucket and then assigned to the instance profile.

Option A and D are invalid, because the credentials should not be provided by a AWS EC2 Instance. Option B is invalid because you would not create a service account, instead you should create an instance profile.

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

- [http://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles\\_use\\_switch-role-ec2-instance-profiles.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2-instance-profiles.html)

**NEW QUESTION 25**

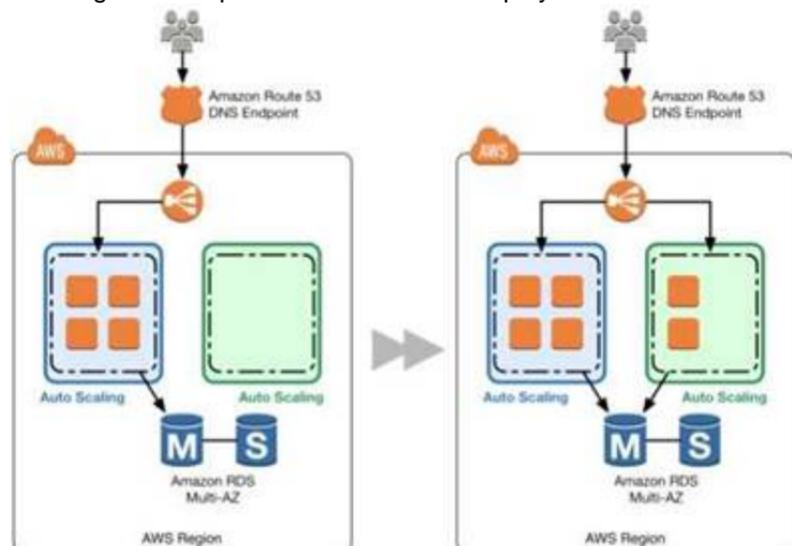
After a daily scrum with your development teams, you've agreed that using Blue/Green style deployments would benefit the team. Which technique should you use to deliver this new requirement?

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

**Answer: C**

**Explanation:**

This is given as a practice in the Green Deployment Guides



A blue group carries the production load while a green group is staged and deployed with the new code. When it's time to deploy, you simply attach the green group to the existing load balancer to introduce traffic to the new environment. For HTTP/HTTPS listeners, the load balancer favors the green Auto Scaling group because it uses a least outstanding requests routing algorithm. As you scale up the green Auto Scaling group, you can take blue Auto Scaling group instances out of service by either terminating them or putting them in Standby state. For more information on Blue Green Deployments, please refer to the below document link: from AWS [https://dOawsstatic.com/whitepapers/AWS\\_Blue\\_Green\\_Deployments.pdf](https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf)

### NEW QUESTION 30

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

- A. Install an Amazon Cloudwatch Logs Agent on every web server during the bootstrap process
- B. Create a CloudWatch log group and define Metric Filters to create custom metrics that track unique visitors from the streaming web server log
- C. Create a scheduled task on an Amazon EC2 instance that runs every hour to generate a new report based on the Cloudwatch custom metric
- D. ^/
- E. On the web servers, create a scheduled task that executes a script to rotate and transmit the logs to Amazon Glacier
- F. Ensure that the operating system shutdown procedure triggers a logs transmission when the Amazon EC2 instance is stopped/terminated
- G. Use Amazon Data Pipeline to process the data in Amazon Glacier and run reports every hour.
- H. On the web servers, create a scheduled task that executes a script to rotate and transmit the logs to an Amazon S3 bucket
- I. Ensure that the operating system shutdown procedure triggers a logs transmission when the Amazon EC2 instance is stopped/terminated
- J. Use AWS Data Pipeline to move log data from the Amazon S3 bucket to Amazon Redshift in order to process and run reports every hour.
- K. Install an AWS Data Pipeline Logs Agent on every web server during the bootstrap process
- L. Create a log group object in AWS Data Pipeline, and define Metric Filters to move processed log data directly from the web servers to Amazon Redshift and run reports every hour.

**Answer:** AC

#### Explanation:

You can use the Cloud Watch Logs agent installer on an existing EC2 instance to install and configure the Cloud Watch Logs agent.

For more information, please visit the below link:

- <http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/QuickStartEC2Instance.html>

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

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

### NEW QUESTION 34

You have a web application that's developed in Node.js. The code is hosted in a Git repository. You want to now deploy this application to AWS. Which of the below 2 options can fulfill this requirement.

- A. Create an Elastic Beanstalk application
- B. Create a Dockerfile to install Node.js
- C. Get the code from Git
- D. Use the command "aws git.push" to deploy the application
- E. Create an AWS CloudFormation template which creates an instance with the AWS::EC2::Container resource type
- F. With UserData, install Git to download the Node.js application and then set it up.
- G. Create a Dockerfile to install Node.js
- H. and get the code from Git
- I. Use the Dockerfile to perform the deployment on a new AWS Elastic Beanstalk application
- J. S
- K. Create an AWS CloudFormation template which creates an instance with the AWS::EC2::Instance resource type and an AMI with Docker pre-installed
- L. With UserData, install Git to download the Node.js application and then set it up.

**Answer:** CD

#### Explanation:

Option A is invalid because there is no "aws git.push" command

Option B is invalid because there is no AWS::EC2::Container resource type.

Elastic Beanstalk supports the deployment of web applications from Docker containers. With Docker containers, you can define your own runtime environment. You can choose your own platform, programming language, and any application dependencies (such as package managers or tools), that aren't supported by other platforms. Docker containers are self-contained and include all the configuration information and software your web application requires to run.

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

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

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 EC2 User data please refer to the below link:

- <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/user-data.html>

Note: "git aws.push" with CB CLI 2.x - see a forum thread at <https://forums.aws.amazon.com/thread.jspa?threadID=583202#jive-message-582979>. Basically, this is a predecessor to the newer "eb deploy" command in CB CLI 3.1. This question kept in order to be consistent with exam.

### NEW QUESTION 35

As an architect you have decided to use CloudFormation instead of OpsWorks or Elastic Beanstalk for deploying the applications in your company. Unfortunately, you have discovered that there is a

resource type that is not supported by CloudFormation. What can you do to get around this.

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

**Answer: B**

**Explanation:**

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

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 39**

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

- A. Sign into the AWS Management Console, copy the old launch configuration, and create a new launch configuration that specifies the C3 instance
- B. Update the Auto Scalinggroup with the new launch configuratio
- C. Auto Scaling will then update the instance type of all running instances.
- D. Sign into the AWS Management Console, and update the existing launch configuration with the new C3 instance typ
- E. Add an UpdatePolicy attribute to your Auto Scaling group that specifies AutoScalingRollingUpdate.
- F. Update the launch configuration specified in the AWS CloudFormation template with the new C3 instance typ
- G. Run a stack update with the new templat
- H. Auto Scaling will then update the instances with the new instance type.
- I. Update the launch configuration specified in the AWS CloudFormation template with the new C3instance typ
- J. Also add an UpdatePolicy attribute to your Auto Scalinggroup that specifies AutoScalingRollingUpdat
- K. Run a stack update with the new template.

**Answer: D**

**Explanation:**

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

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

**NEW QUESTION 41**

What is web identity federation?

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

**Answer: B**

**Explanation:**

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

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

**NEW QUESTION 42**

You need to create a Route53 record automatically in CloudFormation when not running in production during all launches of a Template. How should you implement this?

- A. Use a Parameter for environment, and add a Condition on the Route53 Resource in the template to create the record only when environment is not production.
- B. Create two templates, one with the Route53 record value and one with a null value for the recor
- C. Use the one without it when deploying to production.
- D. Use a Parameter for environment, and add a Condition on the Route53 Resource in the template to create the record with a null string when environment is production.
- E. Create two templates, one with the Route53 record and one without i
- F. Use the one without it when deploying to production.

**Answer: A**

**Explanation:**

The optional Conditions section includes statements that define when a resource is created or when a property is defined. For example, you can compare whether

a value is equal to another value. Based on the result of that condition, you can conditionally create resources. If you have multiple conditions, separate them with commas.

You might use conditions when you want to reuse a template that can create resources in different contexts, such as a test environment versus a production environment. In your template, you can add an EnvironmentType input parameter, which accepts either prod or test as inputs. For the production environment, you might include Amazon EC2 instances with certain capabilities; however, for the test environment, you want to use reduced capabilities to save money. With conditions, you can define which resources are created and how they're configured for each environment type.

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

**NEW QUESTION 46**

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

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

**Answer: B**

**Explanation:**

AWS OpsWorks is a configuration management service that uses Chef, an automation platform that treats server configurations as code. OpsWorks uses Chef to automate how servers are configured, deployed, and managed across your Amazon Elastic Compute Cloud (Amazon EC2) instances or on-premises compute environments. OpsWorks has two offerings, AWS OpsWorks for Chef Automate, and AWS OpsWorks Stacks.

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

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

**NEW QUESTION 51**

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 54**

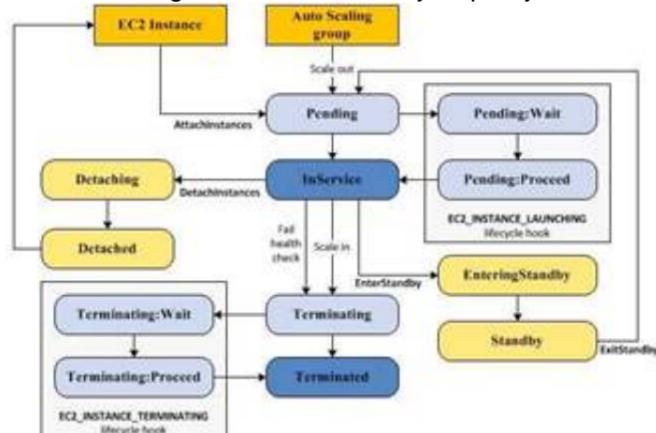
For AWS Auto Scaling, what is the first transition state an existing instance enters after leaving Standby state?

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

**Answer: C**

**Explanation:**

The below diagram shows the Lifecycle policy. When the stand-by state is exited, the next state is pending.



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

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

**NEW QUESTION 58**

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 60**

You need to scale an RDS deployment. You are operating at 10% writes and 90% reads, based on your logging. How best can you scale this in a simple way?

- A. Create a second master RDS instance and peer the RDS groups.
- B. Cache all the database responses on the read side with CloudFront.
- C. Create read replicas for RDS since the load is mostly reads.
- D. Create a Multi-AZ RDS installs and route read traffic to standby.

**Answer:** C

**Explanation:**

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

Option A is invalid because you would need to maintain the synchronization yourself with a secondary instance.

Option B is invalid because you are introducing another layer unnecessarily when you already have read replica's Option D is invalid because you only use this for Standby's

For more information on Read Replica's, please refer to the below link: <https://aws.amazon.com/rds/details/read-replicas/>

**NEW QUESTION 63**

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

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

**Answer:** B

**Explanation:**

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

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

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

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

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

**NEW QUESTION 65**

You have an asynchronous processing application using an Auto Scaling Group and an SQS Queue. The Auto Scaling Group scales according to the depth of the job queue. The completion velocity of the jobs has gone down, the Auto Scaling Group size has maxed out, but the inbound job velocity did not increase. What is a possible issue?

- A. Some of the new jobs coming in are malformed and unprocessable.
- B. The routing tables changed and none of the workers can process events anymore.
- C. Someone changed the IAM Role Policy on the instances in the worker group and broke permissions to access the queue.
- D. The scaling metric is not functioning correctly.

**Answer:** A

**Explanation:**

This question is more on the grounds of validating each option

Option B is invalid, because the Route table would have an effect on all worker processes and no jobs would have been completed.

Option C is invalid because if the IAM Role was invalid then no jobs would be completed.

Option D is invalid because the scaling is happening, its just that the jobs are not getting completed. For more information on Scaling on Demand, please visit the below URL:

- <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scale-based-on-demand.html>

**NEW QUESTION 70**

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

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

**Answer: B**

**Explanation:**

Updated based on latest AWS updates

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

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

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

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

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

**NEW QUESTION 72**

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

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

**Answer: A**

**Explanation:**

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

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

**NEW QUESTION 73**

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

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

**Answer: A**

**Explanation:**

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

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

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

**NEW QUESTION 76**

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

- A. Use the AWS CloudFormation ValidateTemplate call before publishing changes to AWS.
- B. Model your stack in one template, so you can leverage CloudFormation's state management and dependency resolution to propagate all changes.

- C. Use CloudFormation to create brand new infrastructure for all stateless resources on each push, and run integration tests on that set of infrastructure.  
D. Parametrize the template and use Mappings to ensure your template works in multiple Regions.

**Answer:** B

**Explanation:**

Answer - B

Some of the best practices for Cloudformation are

- Created Nested stacks

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

- Reuse Templates

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

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

**NEW QUESTION 79**

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

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

**Answer:** D

**Explanation:**

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

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

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

**NEW QUESTION 82**

You need to grant a vendor access to your AWS account. They need to be able to read protected messages in a private S3 bucket at their leisure. They also use AWS. What is the best way to accomplish this?

- A. Create an IAM User with API Access Key  
B. Grant the User permissions to access the bucket  
C. Give the vendor the AWS Access Key ID and AWS Secret Access Key for the User.  
D. Create an EC2 Instance Profile on your account  
E. Grant the associated IAM role full access to the bucket  
F. Start an EC2 instance with this Profile and give SSH access to the instance to the vendor.  
G. Create a cross-account IAM Role with permission to access the bucket, and grant permission to use the Role to the vendor AWS account.  
D- Generate a signed S3 PUT URL and a signed S3 GET URL, both with wildcard values and 2 year duration  
H. Pass the URLs to the vendor.

**Answer:** C

**Explanation:**

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

- <https://aws.amazon.com/blogs/security/tag/cross-account-access/>

**NEW QUESTION 83**

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

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

**Answer:** B

**Explanation:**

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

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

returns  
an empty list.

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

**NEW QUESTION 87**

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

- A. Use Cloudtrail to log all the API calls, and then traverse the log files to locate the issue
- B. Use Cloudwatch, monitor the CPU utilization to see the times when the CPU peaked
- C. Review the Elastic Load Balancer logs
- D. Create some custom Cloudwatch metrics which are pertinent to the key features of your application

**Answer:** D

**Explanation:**

Since the issue is occurring after the new feature has been added, it could be relevant to the new feature.

Enabling Cloudtrail will just monitor all the API calls of all services and will not benefit the cause.

The monitoring of CPU utilization will just reverify that there is an issue but will not help pinpoint the issue.

The Elastic Load Balancer logs will also just reverify that there is an issue but will not help pinpoint the issue.

For more information on custom Cloudwatch metrics, please refer to the below link:

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

**NEW QUESTION 92**

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 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 93**

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

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

**Answer:** AB

**Explanation:**

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

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

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

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

**NEW QUESTION 94**

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

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

**Answer:** C

**Explanation:**

The AWS documentation mentions

You share resources in one account with users in a different account. By setting up cross-account access in this way, you don't need to create individual IAM users in each account. In addition, users don't have to sign out of one account and sign into another in order to access resources that are in different AWS accounts. After configuring the role, you see how to use the role from the AWS Management Console, the AWS CLI, and the API

For more information on Cross Account Roles Access, please refer to the below link:

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

**NEW QUESTION 98**

You currently have an application with an Auto Scaling group with an Elastic Load Balancer configured in AWS. After deployment users are complaining of slow response time for your application. Which of the following can be used as a start to diagnose the issue

- A. Use Cloudwatch to monitor the HealthyHostCount metric
- B. Use Cloudwatch to monitor the ELB latency
- C. Use Cloudwatch to monitor the CPU Utilization
- D. Use Cloudwatch to monitor the Memory Utilization

**Answer: B**

**Explanation:**

High latency on the ELB side can be caused by several factors, such as:

- Network connectivity
- ELB configuration
- Backend web application server issues

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

- <https://aws.amazon.com/premiumsupport/knowledge-center/elb-latency-troubleshooting/>

**NEW QUESTION 102**

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

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

**Answer: B**

**Explanation:**

The AWS documentation mentions

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

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

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

**NEW QUESTION 105**

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

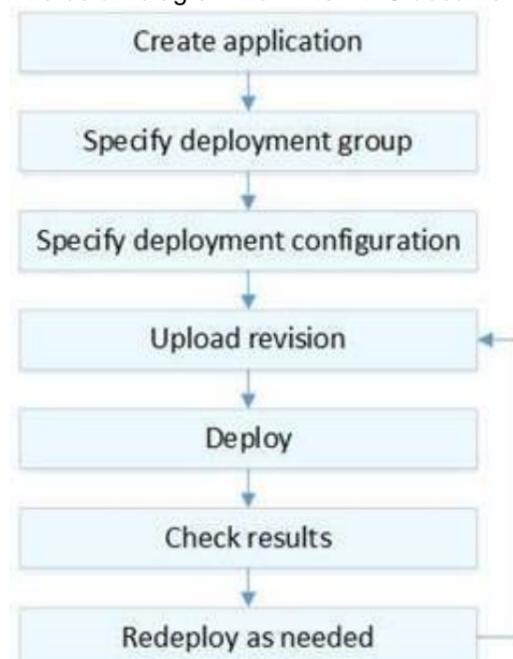
- 1) Specify deployment configuration
- 2) Upload revision
- 3) Create application
- 4) Specify deployment group

- 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 110**

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

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

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 113**

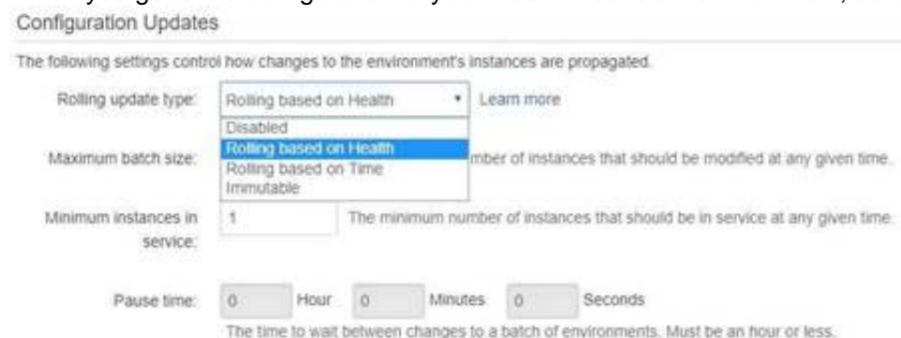
Which of the following is not a rolling type update which is present for Configuration Updates when it comes to the Elastic Beanstalk service

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

**Answer: B**

**Explanation:**

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



The AWS Documentation mentions

- 1) With health-based rolling updates. Elastic Beanstalk waits until instances in a batch pass health checks before moving on to the next batch.
  - 2) For time-based rolling updates, you can configure the amount of time that Elastic Beanstalk waits after completing the launch of a batch of instances before moving on to the next batch. This pause time allows your application to bootstrap and start serving requests.
  - 3) Immutable environment updates are an alternative to rolling updates that ensure that configuration changes that require replacing instances are applied efficiently and safely. If an immutable environment update fails, the rollback process requires only terminating an Auto Scalinggroup. A failed rolling update, on the other hand, requires performing an additional rolling update to roll back the changes.
- For more information on Rolling updates for Elastic beanstalk configuration updates, please visit the below URL:  
 • <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.ro11ingupdates.html>

**NEW QUESTION 118**

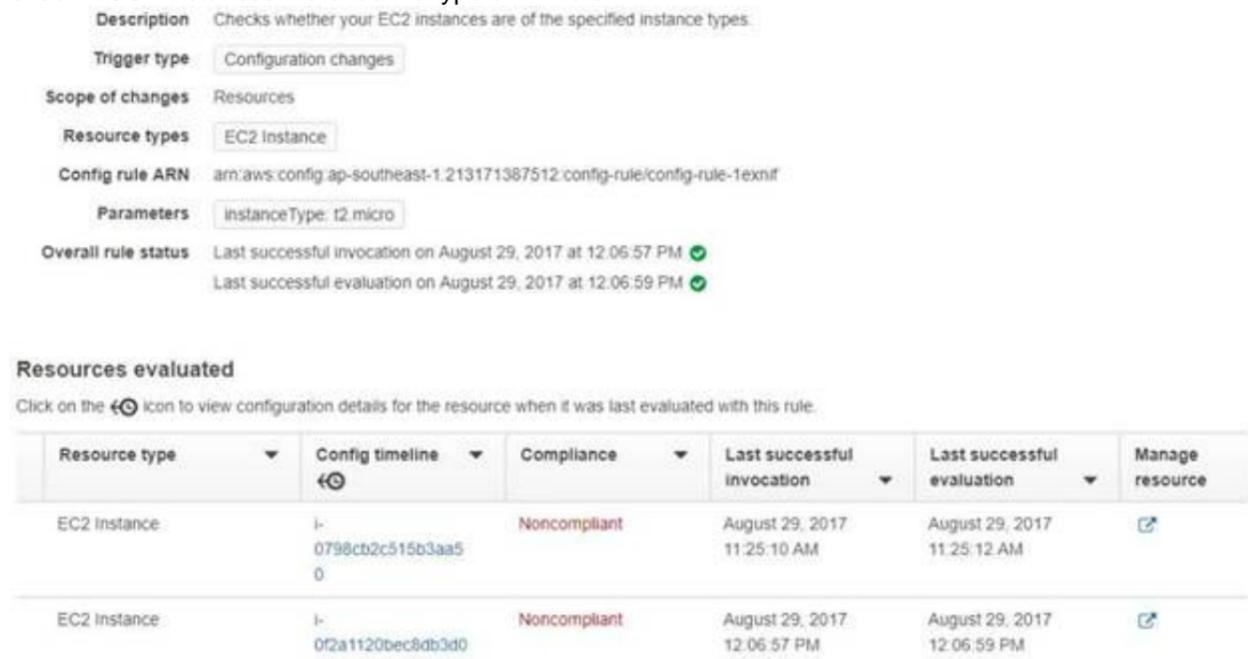
Your company has a set of EC2 resources hosted on AWS. Your new IT procedures state that AWS EC2 Instances must be of a particular Instance type. Which of the following can be used to get the list of EC2 Instances which currently don't match the instance type specified in the new IT procedures

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

**Answer: B**

**Explanation:**

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



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

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

**NEW QUESTION 120**

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

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

**Answer: C**

**Explanation:**

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

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

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

**NEW QUESTION 122**

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 125**

You are in charge 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 128**

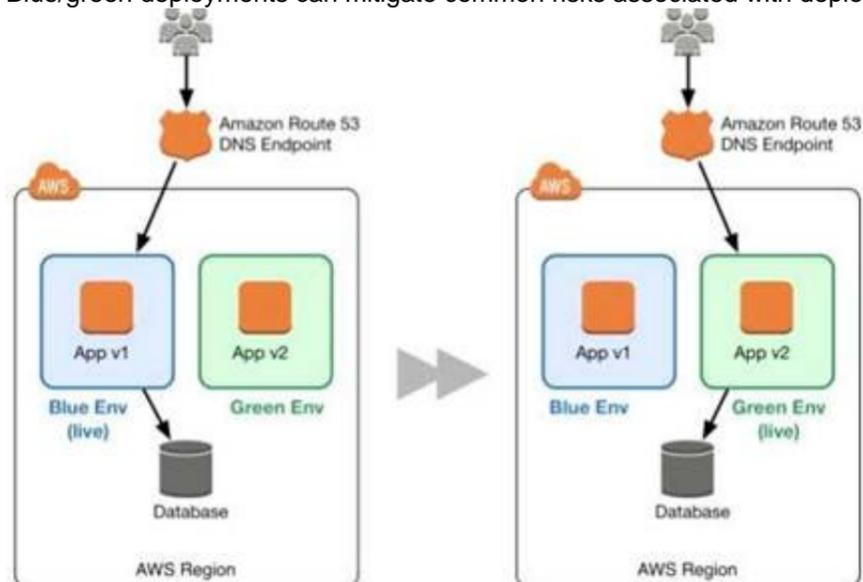
A company has developed a Ruby on Rails content management platform. Currently, OpsWorks with several stacks for dev, staging, and production is being used to deploy and manage the application. Now the company wants to start using Python instead of Ruby. How should the company manage the new deployment?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

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

**Answer: B**

**Explanation:**

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



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

- [https://d03wsstatic.com/whitepapers/AWS\\_Blue\\_Green\\_Deployments.pdf](https://d03wsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf)

**NEW QUESTION 133**

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

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

**Answer:** A

**Explanation:**

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

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

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

**NEW QUESTION 134**

You are in charge of designing CloudFormation templates for your company. One of the key requirements is to ensure that if a CloudFormation stack is deleted, a snapshot of the relational database is created which is part of the stack. How can you achieve this in the best possible way?

- A. Create a snapshot of the relational database beforehand so that when the CloudFormation stack is deleted, the snapshot of the database will be present.
- B. Use the Update policy of the CloudFormation template to ensure a snapshot is created of the relational database.
- C. Use the Deletion policy of the CloudFormation template to ensure a snapshot is created of the relational database.
- D. Create a new CloudFormation template to create a snapshot of the relational database.

**Answer:** C

**Explanation:**

The AWS documentation mentions the following

With the Deletion Policy attribute you can preserve or (in some cases) backup a resource when its stack is deleted. You specify a DeletionPolicy attribute for each resource that you want to control. If a resource has no DeletionPolicy attribute, AWS CloudFormation deletes the resource by default. Note that this capability also applies to update operations that lead to resources being removed.

For more information on the Deletion policy, please visit the below URL: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html>

**NEW QUESTION 135**

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

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

D. Add an RDS MySQL read replica in each availability zone

**Answer:** AD

**Explanation:**

Implement Read Replicas and ElastiCache

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

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

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

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

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

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

**NEW QUESTION 137**

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 139**

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

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

**Answer:** C

**Explanation:**

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

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

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

**NEW QUESTION 143**

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 endpoint 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 (CNAMR) 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 147

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

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

**Answer:** B

#### Explanation:

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

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

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

#### NEW QUESTION 149

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

- A. UpdateCustomCookbooks
- B. ExecuteRecipes
- C. Configure
- D. UnDeploy

**Answer:** ABC

#### NEW QUESTION 151

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

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

**Answer:** AD

#### Explanation:

Implement Read Replicas and Elastic Cache

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

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

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

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

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

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

#### NEW QUESTION 155

Your company is using an Auto Scaling Group to scale out and scale in instances. There is an expectation of a peak in traffic every Monday at 8am. The traffic is then expected to come down before the weekend on Friday 5pm. How should you configure Auto Scaling in this?

- A. Create dynamic scaling policies to scale up on Monday and scale down on Friday
- B. Create a scheduled policy to scale up on Friday and scale down on Monday
- C. Create a scheduled policy to scale up on Monday and scale down on Friday
- D. Manually add instances to the Auto Scaling Group on Monday and remove them on Friday

**Answer:** C

#### Explanation:

The AWS Documentation mentions the following for Scheduled scaling

Scaling based on a schedule allows you to scale your application in response to predictable load changes. For example, every week the traffic to your web application starts to increase on Wednesday, remains high on Thursday, and starts to decrease on Friday. You can plan your scaling activities based on the predictable traffic patterns of your web application.

For more information on scheduled scaling for Auto Scaling, please visit the below URL

• [http://docs.aws.amazon.com/autoscaling/latest/userguide/schedule\\_time.html](http://docs.aws.amazon.com/autoscaling/latest/userguide/schedule_time.html)

**NEW QUESTION 157**

Your company has recently extended its datacenter into a VPC on AWS. There is a requirement for on-premise users manage AWS resources from the AWS console. You don't want to create IAM users for them again. Which of the below options will fit your needs for authentication?

- A. Use OAuth 2.0 to retrieve temporary AWS security credentials to enable your members to sign in to the AWS Management Console.
- B. Use web Identity Federation to retrieve AWS temporary security credentials to enable your members to sign in to the AWS Management Console.
- C. Use your on-premises SAML 2 O-compliant identity provider (IDP) to grant members federated access to the AWS Management Console via the AWS single sign-on (SSO) endpoint.
- D. Use your on-premises SAML 2.0-compliant identity provider (IDP) to retrieve temporary security credentials to enable members to sign in to the AWS Management Console.

**Answer: C**

**Explanation:**

You can use a role to configure your SAML 2.0-compliant IdP and AWS to permit your federated users to access the AWS Management Console. The role grants the user permissions to carry out tasks in the console.

For more information on aws SAML, please visit the below URL

• [http://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles\\_providers\\_enable-console-saml.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_enable-console-saml.html)

**NEW QUESTION 161**

Explain what the following resource in a CloudFormation template does? Choose the best possible answer.

```
"SNSTopic" : {
  "Type" : "AWS::SNS::Topic",
  "Properties" : {
    "Subscription" : [ {
      "Protocol" : "sqs",
      "Endpoint" : { "Fn::GetAtt" : [ "SQSQueue", "Arn" ] }
    } ]
  }
}
```

- A. Creates an SNS topic which allows SQS subscription endpoints to be added as a parameter on the template
- B. Creates an SNS topic that allow SQS subscription endpoints
- C. Creates an SNS topic and then invokes the call to create an SQS queue with a logical resource name of SQSQueue
- D. Creates an SNS topic and adds a subscription ARN endpoint for the SQS resource created under the logical name SQSQueue

**Answer: D**

**Explanation:**

The intrinsic function Fn::GetAtt returns the value of an attribute from a resource in the template. This has nothing to do with adding parameters (Option A is wrong) or allowing endpoints (Option B is wrong) or invoking relevant calls (Option C is wrong)

For more information on Fn::GetAtt function please refer to the below link

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

**NEW QUESTION 164**

Which of the following is not a component of Elastic Beanstalk?

- A. Application
- B. Environment
- C. Docker
- D. ApplicationVersion

**Answer: C**

**Explanation:**

Answer - C

The following are the components of Elastic Beanstalk

- 1) Application - An Elastic Beanstalk application is a logical collection of Elastic Beanstalk components, including environments, versions, and environment configurations. In Elastic Beanstalk an application is conceptually similar to a folder
- 2) Application version - In Elastic Beanstalk, an application version refers to a specific, labeled iteration of deployable code for a web application
- 3) environment - An environment is a version that is deployed onto AWS resources. Each environment runs only a single application version at a time, however you can run the same version or different versions in many environments at the same time.
- 4) environment Configuration - An environment configuration identifies a collection of parameters and settings that define how an environment and its associated resources behave.
- 5) Configuration Template - A configuration template is a starting point for creating unique environment configurations. For more information on the components of Elastic beanstalk please refer to the below link

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

**NEW QUESTION 166**

If you're trying to configure an AWS Elastic Beanstalk worker tier for easy debugging if there are problems finishing queue jobs, what should you configure?

- A. ConfigureRolling Deployments.
- B. ConfigureEnhanced Health Reporting.
- C. ConfigureBlue-Green Deployments.
- D. Configure a Dead Letter Queue.

**Answer:** D

**Explanation:**

The AWS documentation mentions the following on dead-letter queues

Amazon SQS supports dead-letter queues. A dead-letter queue is a queue that other (source) queues can target for messages that can't be processed (consumed)

successfully. You can set aside and isolate these messages in the dead-letter queue to determine why their processing doesn't succeed.

For more information on dead letter queues please visit the below link <http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-dead-letter-queues.html>

**NEW QUESTION 167**

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 172**

You are in charge of designing a CloudFormation template which deploys a LAMP stack. After deploying a stack, you see that the status of the stack is showing as CREATE\_COMPLETE, but the Apache server is still not up and running and is experiencing issues while starting up. You want to ensure that the stack creation only shows the status of CREATE\_COMPLETE after all resources defined in the stack are up and running. How can you achieve this?

Choose 2 answers from the options given below.

- A. Define a stack policy which defines that all underlying resources should be up and running before showing a status of CREATE\_COMPLETE.
- B. Use lifecycle hooks to mark the completion of the creation and configuration of the underlying resource.
- C. Use the CreationPolicy to ensure it is associated with the EC2 Instance resource.
- D. Use the CFN helper scripts to signal once the resource configuration is complete.

**Answer:** CD

**Explanation:**

The AWS Documentation mentions

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

For more information on the CreationPolicy, please visit the below URL <https://aws.amazon.com/blogs/devops/use-a-creationpolicy-to-wait-for-on-instance-configurations/>

**NEW QUESTION 174**

Which of the following are true with regard to OpsWorks stack Instances? Choose 3 answers from the options given below.

- A. A stack's instances can be a combination of both Linux and Windows based operating systems.
- B. You can use EC2 Instances that were created outside the boundary of OpsWorks.
- C. You can use instances running on your own hardware.
- D. You can start and stop instances manually.

**Answer:** BCD

**Explanation:**

The AWS Documentation mentions the following

1) You can start and stop instances manually or have AWS OpsWorks Stacks automatically scale the number of instances. You can use time-based automatic scaling with any stack; Linux stacks also can use load-based scaling.

2) In addition to using AWS OpsWorks Stacks to create Amazon Linux instances, you can also register instances with a Linux stack that were created outside of AWS OpsWorks Stacks. This includes Amazon EC2 instances and instances running on your own hardware. However, they must be running one of the supported Linux distributions. You cannot register Amazon EC2 or on-premises Windows instances.

3) A stack's instances can run either Linux or Windows. A stack can have different Linux versions or distributions on different instances, but you cannot mix Linux and Windows instances.

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

**NEW QUESTION 177**

Which of the following Cloudformation helper scripts can help install packages on EC2 resources

- A. cfn-init
- B. cfn-signal
- C. cfn-get-metadata
- D. cfn-hup

**Answer: A**

**Explanation:**

The AWS Documentation mentions

Currently, AWS CloudFormation provides the following helpers:

cfn-init: Used to retrieve and interpret the resource metadata, installing packages, creating files and starting services.

cfn-signal: A simple wrapper to signal an AWS CloudFormation CreationPolicy or WaitCondition, enabling you to synchronize other resources in the stack with the application being ready.

cfn-get-metadata: A wrapper script making it easy to retrieve either all metadata defined for a resource or path to a specific key or subtree of the resource metadata.

cfn-hup: A daemon to check for updates to metadata and execute custom hooks when the changes are detected. For more information on helper scripts, please visit the below URL: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-helper-scripts-reference.html>

**NEW QUESTION 180**

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

**Continuous Integration**



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

**NEW QUESTION 184**

Which of the following services along with Cloudformation helps in building a Continuous Delivery release practice

- A. AWSConfig
- B. AWSLambda
- C. AWSCloudtrail
- D. AWSCodePipeline

**Answer: D**

**Explanation:**

The AWS Documentation mentions

Continuous delivery is a release practice in which code changes are automatically built, tested, and prepared for release to production. With AWS Cloud Formation and AWS CodePipeline, you can use continuous delivery to automatically build and test changes to your AWS Cloud Formation templates before promoting them to

production stacks. This release process lets you rapidly and reliably make changes to your AWS infrastructure.

For more information on Continuous Delivery, please visit the below URL: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/continuous-delivery-codepipeline.html>

**NEW QUESTION 187**

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. Autoscalingonly works for single availability zone
- B. Autoscalingcan be enabled for multi AZ only in north Virginia region
- C. Availabilityzone is not added to Elastic load balancer
- D. Instancesneed to manually added to availability zone

**Answer: C**

**Explanation:**

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

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

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

**NEW QUESTION 188**

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

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

**Answer:** C

**Explanation:**

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

For more information on IAM please see the below link:

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

**NEW QUESTION 190**

What are the benefits when you implement a Blue Green deployment for your infrastructure or application level changes. Choose 3 answers from the options given below

- A. Near zero-downtime release for new changes
- B. Better rollback capabilities
- C. Ability to deploy with higher risk
- D. Good turnaround time for application deployments

**Answer:** ABD

**Explanation:**

The AWS Documentation mentions the following

Blue/green deployments provide near zero-downtime release and rollback capabilities. The fundamental idea behind blue/green deployment is to shift traffic between two identical environments that are running different versions of your application. The blue environment represents the current application version serving production traffic. In parallel, the green environment is staged running a different version of your application. After the green environment is ready and tested, production traffic is redirected from blue to green.

For more information on Blue Green deployments please see the below link:

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

**NEW QUESTION 191**

Your team is responsible for an AWS Elastic Beanstalk application. The business requires that you move to a continuous deployment model, releasing updates to the application multiple times per day with zero downtime. What should you do to enable this and still be able to roll back almost immediately in an emergency to the previous version?

- A. Enable rolling updates in the Elastic Beanstalk environment, setting an appropriate pause time for application startup.
- B. Create a second Elastic Beanstalk environment running the new application version, and swap the environment CNAMEs.
- C. Develop the application to poll for a new application version in your code repository; download and install to each running Elastic Beanstalk instance.
- D. Create a second Elastic Beanstalk environment with the new application version, and configure the old environment to redirect clients, using the HTTP 301 response code, to the new environment

**Answer:** B

**Explanation:**

The AWS Documentation mentions the below

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 CNAMEs of the two environments to redirect traffic to the new version instantly. For more information on Elastic Beanstalk swap URL please see the below link:

- [http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAME\\_CSwap.html](http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAME_CSwap.html)

**NEW QUESTION 195**

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

- A. Publish your data to CloudWatch Logs, and configure your application to autoscale to handle the load on demand.
- B. Publish your log data to an Amazon S3 bucket
- C. Use AWS CloudFormation to create an AutoScaling group to scale your post-processing application which is configured to pull down your log files stored in Amazon S3.
- D. Post your log data to an Amazon Kinesis data stream, and subscribe your log-processing application so that it is configured to process your logging data.

E. Configure an Auto Scaling group to increase the size of your Amazon EMR cluster

**Answer: C**

**Explanation:**

The AWS Documentation mentions the below

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

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

Amazon Kinesis enables you to process and analyze data as it arrives and respond in real-time instead of having to wait until all your data is collected before the processing can begin.

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

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

**NEW QUESTION 197**

You are managing the development of an application that uses DynamoDB to store JSON data. You have already set the Read and Write capacity of the DynamoDB table. You are unsure of the amount of the traffic that will be received by the application during the deployment time. How can you ensure that the DynamoDB is not highly throttled and does not become a bottleneck for the application? Choose 2 answers from the options below.

- A. Monitor the ConsumedReadCapacityUnits and ConsumedWriteCapacityUnits metric using Cloudwatch.
- B. Monitor the SystemErrors metric using Cloudwatch
- C. Create a Cloudwatch alarm which would then send a trigger to AWS Lambda to increase the Read and Write capacity of the DynamoDB table.
- D. Create a Cloudwatch alarm which would then send a trigger to AWS Lambda to create a new DynamoDB table.

**Answer: AC**

**Explanation:**

Refer to the following AWS Documentation that specifies what should be monitored for a DynamoDB table.

How can I determine how much of my provisioned throughput is being used?	You can monitor ConsumedReadCapacityUnits or ConsumedWriteCapacityUnits over the specified time period, to track how much of your provisioned throughput is being used.
--	---

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

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

**NEW QUESTION 198**

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

- A. Amazon EC2 Instances
- B. AWS Lambda Functions
- C. Amazon CodeCommit
- D. Amazon ECS Tasks

**Answer: ABD**

**Explanation:**

The AWS Documentation mentions the below

You can configure the following AWS services as targets for Cloud Watch Events

- Amazon EC2 instances
- AWS Lambda functions
- Streams in Amazon Kinesis Streams
- Delivery streams in Amazon Kinesis Firehose
- Amazon ECS tasks
- SSM Run Command
- SSM Automation
- Step Functions state machines
- Pipelines in AWS CodePipeline
- Amazon Inspector assessment templates
- Amazon SNS topics
- Amazon SQS queues
- Built-in targets
- The default event bus of another AWS account

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

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

**NEW QUESTION 202**

You are planning on configuring logs for your Elastic Load balancer. At what intervals does the logs get produced by the Elastic Load balancer service. Choose 2 answers from the options given below

- A. 5 minutes
- B. 60 minutes

- C. 1 minute
- D. 30seconds

**Answer:** AB

**Explanation:**

The AWS Documentation mentions

Elastic Load Balancing publishes a log file for each load balancer node at the interval you specify. You can specify a publishing interval of either 5 minutes or 60 minutes when you enable the access log for your load balancer. By default, Elastic Load Balancing publishes logs at a 60-minute interval.

For more information on Elastic load balancer logs please see the below link: <http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/access-log-collection.html>

**NEW QUESTION 203**

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. You don't have the necessary access to the logs generated by ELB.
- B. By default ELB access logs are disabled.
- C. The Autoscaling service is not sending the required logs to ELB
- D. The EC2 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, Elastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify. You can disable access logging at any time.

For more information on ELB access logs please see the below link:

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

**NEW QUESTION 207**

Which of the following tools is available to send log data from EC2 Instances.

- A. CloudWatch LogsAgent
- B. CloudWatchAgent
- C. Logsconsole.
- D. LogsStream

**Answer:** A

**Explanation:**

The AWS Documentation mentions the following

The CloudWatch Logs agent provides an automated way to send log data to Cloud Watch Logs from Amazon EC2 instances. The agent is comprised of the following components:

A plug-in to the AWS CLI that pushes log data to CloudWatch Logs.

A script (daemon) that initiates the process to push data to CloudWatch Logs.

Acron job that ensures that the daemon is always running. For more information on Cloudwatch logs Agent, please see the below link:

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

**NEW QUESTION 209**

You have configured the following AWS services in your organization - Auto Scaling group, Elastic Load Balancer, and EC2 instances. You have been requested to terminate an instance from the Auto Scaling Group when the CPU utilization is less than 30%. How can you achieve this.

- A. Create a Cloud Watch alarm to send a notification to SQS
- B. SQS can then remove one instance from the Auto Scaling Group.
- C. Create a CloudWatch alarm to send a notification to the Auto Scaling group when the aggregated CPU utilization is less than 30% and configure the Auto Scaling policy to remove one instance.
- D. Create a CloudWatch alarm to send a notification to the ELB
- E. The ELB can then remove one instance from the Auto Scaling Group.
- F. Create a CloudWatch alarm to send a notification to the admin team
- G. The admin team can then manually terminate an instance from the Auto Scaling Group.

**Answer:** B

**Explanation:**

The AWS Documentation mentions the following

You should have two policies, one for scaling in (terminating instances) and one for scaling out (launching instances), for each event to monitor. For example, if you want to scale out when the network bandwidth reaches a certain level, create a policy specifying that Auto Scaling should start a certain number of instances to help with your traffic. But you may also want an accompanying policy to scale in by a certain number when the network bandwidth level goes back down

For more information on the scaling plans, please see the below link: [http://docs.aws.amazon.com/autoscaling/latest/userguide/scaling\\_plan.html](http://docs.aws.amazon.com/autoscaling/latest/userguide/scaling_plan.html)

**NEW QUESTION 212**

Which of the following is a reliable and durable logging solution to track changes made to your AWS resources?

- A. Create a new CloudTrail trail with one new S3 bucket to store the logs and with the global services option selected
- B. Use IAM roles, S3 bucket policies and MultiFactor Authentication (MFA) Delete on the S3 bucket that stores your log
- C. V
- D. Create a new CloudTrail with one new S3 bucket to store the log
- E. Configure SNS to send log file delivery notifications to your management system

- F. Use 1AM roles and S3 bucket policies on the S3 bucket that stores your logs.
- G. Create a new CloudTrail trail with an existing S3 bucket to store the logs and with the global services option selected.
- H. Use S3 ACLs and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs.
- I. Create three new CloudTrail trails with three new S3 buckets to store the logs one for the AWS Management console, one for AWS SDKs and one for command line tools. Use 1AM roles and S3 bucket policies on the S3 buckets that store your logs.

**Answer:** A

**Explanation:**

AWS Identity and Access Management (IAM) is integrated with AWS CloudTrail, a service that logs AWS events made by or on behalf of your AWS account. CloudTrail logs authenticated AWS API calls and also AWS sign-in events, and collects this event information in files that are delivered to Amazon S3 buckets. You need to ensure that all services are included. Hence option B is partially correct.

Option B and D is wrong because it just adds an overhead for having 3 S3 buckets and SNS notifications.

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

- <http://docs.aws.amazon.com/IAM/latest/UserGuide/cloudtrail-integration.html>

**NEW QUESTION 217**

By default in Opsworks, how many application versions can you rollback up to?

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

**Answer:** D

**Explanation:**

The AWS Documentation mentions the following Restores the previously deployed app version. For example, if you have deployed the app three times and then run Rollback, the server will serve the app from the second deployment. If you run Rollback again, the server will serve the app from the first deployment. By default, AWS OpsWorks Stacks stores the five most recent deployments, which allows you to roll back up to four versions. If you exceed the number of stored versions, the command fails and leaves the oldest version in place.

For more information on Opsworks app deployment, please visit the below URL: <http://docs.aws.amazon.com/opsworks/latest/userguide/workingapps-deploying.html>

**NEW QUESTION 222**

Your application requires long-term storage for backups and other data that you need to keep readily available but with lower cost. Which S3 storage option should you use?

- A. AmazonS3 Standard- Infrequent Access  
B. S3Standard  
C. Glacier  
D. ReducedRedundancy Storage

**Answer:** A

**Explanation:**

The AWS Documentation mentions the following

Amazon S3 Standard - Infrequent Access (Standard - IA) is an Amazon S3 storage class for data that is accessed less frequently, but requires rapid access when needed. Standard - IA offers the high durability, throughput, and low latency of Amazon S3 Standard, with a low per GB storage price and per GB retrieval fee.

For more information on S3 Storage classes, please visit the below URL:

- <https://aws.amazon.com/s3/storage-classes/>

**NEW QUESTION 224**

You are designing an application that contains protected health information. Security and compliance requirements for your application mandate that all protected health information in the application use encryption at rest and in transit. The application uses a three-tier architecture where data flows through the load balancer and is stored on Amazon EBS volumes for processing and the results are stored in Amazon S3 using the AWS SDK.

Which of the following two options satisfy the security requirements? (Select two)

- A. Use SSL termination on the load balancer, Amazon EBS encryption on Amazon EC2 instances and Amazon S3 with server-side encryption.
- B. Use SSL termination with a SAN SSL certificate on the load balancer
- C. Amazon EC2 with all Amazon EBS volumes using Amazon EBS encryption, and Amazon S3 with server-side encryption with customer-managed keys.
- D. Use TCP load balancing on the load balancer
- E. SSL termination on the Amazon EC2 instance
- F. OS-level disk encryption on the Amazon EBS volumes and Amazon S3 with server-side encryption.
- G. Use TCP load balancing on the load balancer
- H. SSL termination on the Amazon EC2 instances and Amazon S3 with server-side encryption.
- I. Use SSL termination on the load balancer, an SSL listener on the Amazon EC2 instances, Amazon EBS encryption on EBS volumes containing PHI and Amazon S3 with server-side encryption.

**Answer:** CE

**Explanation:**

The AWS Documentation mentions the following: HTTPS/SSL Listeners

You can create a load balancer with the following security features. SSL Server Certificates

If you use HTTPS or SSL for your front-end connections, you must deploy an X.509 certificate (SSL server certificate) on your load balancer. The load balancer decrypts

requests from clients before sending them to the back-end instances (known as SSL termination). For more information, see SSL/TLS Certificates for Classic Load Balancers.

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.

Reference Link:

? <http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-listener-config.html>

Create a Classic Load Balancer with an HTTPS Listener

A load balancer takes requests from clients and distributes them across the EC2 instances that are registered with the load balancer.

You can create a load balancer that listens on both the HTTP (80) and HTTPS (443) ports. If you specify that the HTTPS listener sends requests to the instances on port 80, the load balancer terminates the requests and communication from the load balancer to the instances is not encrypted. If the HTTPS listener sends requests to the instances on port 443, communication from the load balancer to the instances is encrypted.

Reference Link:

• <http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-create-https-ssl-load-balancer.html> | Option A & B are incorrect because they are missing encryption in transit between ELB and EC2 instances.

Option D is incorrect because it is missing encryption at rest on the data associated with the EC2 instances.

## NEW QUESTION 225

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

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

\* One year free update

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

\* Trusted by Millions

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

\* Shop Securely

All transactions are protected by VeriSign!

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

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