



**Amazon**

## **Exam Questions AWS-Certified-DevOps-Engineer-Professional**

Amazon AWS Certified DevOps Engineer Professional

### NEW QUESTION 1

Your application consists of 10% writes and 90% reads. You currently service all requests through a Route53 Alias Record directed towards an AWS ELB, which sits in front of an EC2 Auto Scaling Group. Your system is getting very expensive when there are large traffic spikes during certain news events, during which many more people request to read similar data all at the same time. What is the simplest and cheapest way to reduce costs and scale with spikes like this?

- A. Create an S3 bucket and asynchronously replicate common requests responses into S3 object
- B. When a request comes in for a precomputed response, redirect to AWS S3.
- C. Create another ELB and Auto Scaling Group layer mounted on top of the other system, adding a tier to the system
- D. Serve most read requests out of the top layer.
- E. Create a CloudFront Distribution and direct Route53 to the Distribution
- F. Use the ELB as an Origin and specify Cache Behaviours to proxy cache requests which can be served late.
- G. Create a Memcached cluster in AWS ElastiCache
- H. Create cache logic to serve requests which can be served late from the in-memory cache for increased performance.

**Answer: C**

#### Explanation:

CloudFront is ideal for scenarios in which entire requests can be served out of a cache and usage patterns involve heavy reads and spikiness in demand. A cache behavior is the set of rules you configure for a given URL pattern based on file extensions, file names, or any portion of a URL path on your website (e.g., \*.jpg). You can configure multiple cache behaviors for your web distribution. Amazon CloudFront will match incoming viewer requests with your list of URL patterns, and if there is a match, the service will honor the cache behavior you configure for that URL pattern. Each cache behavior can include the following Amazon CloudFront configuration values: origin server name, viewer connection protocol, minimum expiration period, query string parameters, cookies, and trusted signers for private content.

Reference: <https://aws.amazon.com/Cloudfront/dynamic-content/>

### NEW QUESTION 2

You run a clustered NoSQL database on AWS EC2 using AWS EBS. You need to reduce latency for database response times. Performance is the most important concern, not availability. You did not perform the initial setup, someone without much AWS knowledge did, so you are not sure if they configured everything optimally. Which of the following is NOT likely to be an issue contributing to increased latency?

- A. The EC2 instances are not EBS Optimized.
- B. The database and requesting system are both in the wrong Availability Zone.
- C. The EBS Volumes are not using PIOPS.
- D. The database is not running in a placement group

**Answer: B**

#### Explanation:

For the highest possible performance, all instances in a clustered database like this one should be in a single Availability Zone in a placement group, using EBS optimized instances, and using PIOPS SSD EBS Volumes. The particular Availability Zone the system is running in should not be important, as long as it is the same as the requesting resources.

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

### NEW QUESTION 3

You are creating an application which stores extremely sensitive financial information. All information in the system must be encrypted at rest and in transit. Which of these is a violation of this policy?

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

**Answer: A**

#### Explanation:

Terminating SSL terminates the security of a connection over HTTP, removing the S for "Secure" in HTTPS. This violates the "encryption in transit" requirement in the scenario.

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-listener-config.html>

### NEW QUESTION 4

When thinking of AWS Elastic Beanstalk, the 'Swap Environment URLs' feature most directly aids in what?

- A. Immutable Rolling Deployments
- B. Mutable Rolling Deployments
- C. Canary Deployments
- D. Blue-Green Deployments

**Answer: D**

#### Explanation:

Simply upload the new version of your application and let your deployment service (AWS Elastic Beanstalk, AWS CloudFormation, or AWS OpsWorks) deploy a new version (green). To cut over to the new version, you simply replace the ELB URLs in your DNS records. Elastic Beanstalk has a Swap Environment URLs feature to facilitate a simpler cutover process.

Reference: <https://d0.awsstatic.com/whitepapers/overview-of-deployment-options-on-aws.pdf>

### NEW QUESTION 5

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

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

**Answer: B**

**Explanation:**

Because you only need to batch analyze, anything using streaming is a waste of money. CloudFront is a Gigabit-Scale HTTP(S) global request distribution service, so it can handle scale, geo-spread, spikes, and unpredictability. The Access Logs will contain the GET data and work just fine for batch analysis and email using EMR.

Can I use Amazon CloudFront if I expect usage peaks higher than 10 Gbps or 15,000 RPS? Yes. Complete our request for higher limits here, and we will add more capacity to your account within two business days.

Reference: <https://aws.amazon.com/Cloudfront/faqs/>

**NEW QUESTION 6**

To monitor API calls against our AWS account by different users and entities, we can use to create a history of calls in bulk for later review, and use for reacting to AWS API calls in real-time.

- A. AWS Config; AWS Inspector
- B. AWS CloudTrail; AWS Config
- C. AWS CloudTrail; CloudWatch Events
- D. AWS Config; AWS Lambda

**Answer: C**

**Explanation:**

CloudTrail is a batch API call collection service, CloudWatch Events enables real-time monitoring of calls through the Rules object interface.

Reference: <https://aws.amazon.com/whitepapers/security-at-scale-governance-in-aws/>

**NEW QUESTION 7**

How does Amazon RDS multi Availability Zone model work?

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

**Answer: A**

**Explanation:**

In a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone.

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

**NEW QUESTION 8**

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

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

**Answer: D**

**Explanation:**

This is the complete list of Intrinsic Functions...: Fn::Base64, Fn::And, Fn::Equals, Fn::If, Fn::Not, Fn::Or, Fn::FindInMap, Fn::GetAtt, Fn::GetAZs, Fn::Join, Fn::Select, Ref

Reference:

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

**NEW QUESTION 9**

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

- A. Copy all log files into AWS S3 using a cron job on each instanc
- B. Use an S3 Notification Configuration on the `PutBucket` event and publish events to AWS Lambd
- C. Use the Lambda to analyze logs as soon as they come in and flag issues.
- D. Begin using CloudWatch Logs on every servic
- E. Stream all Log Groups into S3 object
- F. Use AWS EMR clusterjobs to perform ad-hoc MapReduce analysis and write new queries when needed.

- G. Copy all log files into AWS S3 using a cron job on each instance
- H. Use an S3 Notification Configuration on the `PutBucket` event and publish events to AWS Kinesis
- I. Use Apache Spark on AWS EMR to perform at-scale stream processing queries on the log chunks and flag issues.
- J. Begin using CloudWatch Logs on every service
- K. Stream all Log Groups into an AWS Elasticsearch Service Domain running Kibana 4 and perform log analysis on a search cluster.

**Answer:** D

**Explanation:**

The Elasticsearch and Kibana 4 combination is called the ELK Stack, and is designed specifically for real-time, ad-hoc log analysis and aggregation. All other answers introduce extra delay or require pre-defined queries.

Amazon Elasticsearch Service is a managed service that makes it easy to deploy, operate, and scale Elasticsearch in the AWS Cloud. Elasticsearch is a popular open-source search and analytics engine for use cases such as log analytics, real-time application monitoring, and click stream analytics. Reference:

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

**NEW QUESTION 10**

You need to create a simple, holistic check for your system's general availability and uptime. Your system presents itself as an HTTP-speaking API. What is the most simple tool on AWS to achieve this with?

- A. Route53 Health Checks
- B. CloudWatch Health Checks
- C. AWS ELB Health Checks
- D. EC2 Health Checks

**Answer:** A

**Explanation:**

You can create a health check that will run into perpetuity using Route53, in one API call, which will ping your service via HTTP every 10 or 30 seconds.

Amazon Route 53 must be able to establish a TCP connection with the endpoint within four seconds. In addition, the endpoint must respond with an HTTP status code of 200 or greater and less than 400 within two seconds after connecting.

Reference:

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-determining-health-of-endpoints.html>

**NEW QUESTION 10**

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

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

**Answer:** B

**Explanation:**

the bootstrapping process can be slower if you have a complex application or multiple applications to install. Managing a fleet of applications with several build tools and dependencies can be a challenging task during rollouts. Furthermore, your deployment service should be designed to do faster rollouts to take advantage of Auto Scaling.

Reference: <https://d0.awsstatic.com/whitepapers/overview-of-deployment-options-on-aws.pdf>

**NEW QUESTION 15**

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

- A. AWS Elasticsearch Service
- B. AWS RedShift
- C. AWS EMR
- D. AWS DynamoDB

**Answer:** A

**Explanation:**

Amazon Elasticsearch Service (Amazon ES) is a managed service that makes it easy to deploy, operate, and scale Elasticsearch clusters in the AWS cloud. Elasticsearch is a popular open-source search and analytics engine for use cases such as log analytics, real-time application monitoring, and click stream analytics.

Reference:

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

**NEW QUESTION 17**

What is the scope of an EC2 EIP?

- A. Placement Group
- B. Availability Zone
- C. Region
- D. VPC

**Answer:** C

**Explanation:**

An Elastic IP address is tied to a region and can be associated only with an instance in the same region. Reference:  
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/resources.html>

**NEW QUESTION 21**

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

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

**Answer: C**

**Explanation:**

You can put any instance that is in an InService state into a Standby state. This enables you to remove the instance from service, troubleshoot or make changes to it, and then put it back into service. Instances in a Standby state continue to be managed by the Auto Scaling group. However, they are not an active part of your application until you put them back into service.

Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AutoScalingGroupLifecycle.html>

**NEW QUESTION 24**

There are a number of ways to purchase compute capacity on AWS. Which orders the price per compute or memory unit from LOW to HIGH (cheapest to most expensive), on average?

(A) On-Demand (B) Spot (C) Reserved

- A. A, B, C
- B. C, B, A
- C. B, C, A
- D. A, C, B

**Answer: C**

**Explanation:**

Spot instances are usually many, many times cheaper than on-demand prices. Reserved instances, depending on their term and utilization, can yield approximately 33% to 66% cost savings. On-Demand prices are the baseline price and are the most expensive way to purchase EC2 compute time. Reference:  
[https://d0.awsstatic.com/whitepapers/Cost\\_Optimization\\_with\\_AWS.pdf](https://d0.awsstatic.com/whitepapers/Cost_Optimization_with_AWS.pdf)

**NEW QUESTION 25**

You run operations for a company that processes digital wallet payments at a very high volume. One second of downtime, during which you drop payments or are otherwise unavailable, loses you on average USD 100. You balance the financials of the transaction system once per day. Which database setup is best suited to address this business risk?

- A. A multi-AZ RDS deployment with synchronous replication to multiple standbys and read-replicas for fast failover and ACID properties.
- B. A multi-region, multi-master, active-active RDS configuration using database-level ACID design principles with database trigger writes for replication.
- C. A multi-region, multi-master, active-active DynamoDB configuration using application control-level BASE design principles with change-stream write queue buffers for replication.
- D. A multi-AZ DynamoDB setup with changes streamed to S3 via AWS Kinesis, for highly durable storage and BASE properties.

**Answer: C**

**Explanation:**

Only the multi-master, multi-region DynamoDB answer makes sense. Multi-AZ deployments do not provide sufficient availability when a business loses USD 360,000 per hour of unavailability. As RDS does not natively support multi-region, and ACID does not perform well/at all over large distances between regions, only the DynamoDB answer works. Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.CrossRegionRepl.html>

**NEW QUESTION 27**

When thinking of DynamoDB, what are true of Local Secondary Key properties?

- A. Either the partition key or the sort key can be different from the table, but not both.
- B. Only the sort key can be different from the table.
- C. The partition key and sort key can be different from the table.
- D. Only the partition key can be different from the table.

**Answer: B**

**Explanation:**

Global secondary index — an index with a partition key and a sort key that can be different from those on the table. A global secondary index is considered "global" because queries on the index can span all of the data in a table, across all partitions.

Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html>

**NEW QUESTION 29**

For AWS Auto Scaling, what is the first transition state a new instance enters after leaving steady state when scaling out due to increased load?

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

D. Detaching

**Answer: B**

**Explanation:**

When a scale out event occurs, the Auto Scaling group launches the required number of EC2 instances, using its assigned launch configuration. These instances start in the Pending state. If you add a lifecycle hook to your Auto Scaling group, you can perform a custom action here. For more information, see Lifecycle Hooks.

Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AutoScalingGroupLifecycle.html>

**NEW QUESTION 31**

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.
- H. Generate a signed S3 PUT URL and a signed S3 GET URL, both with wildcard values and 2 year duration
- I. Pass the URLs to the vendor.

**Answer: C**

**Explanation:**

When third parties require access to your organization's AWS resources, you can use roles to delegate access to them. For example, a third party might provide a service for managing your AWS resources. With IAM roles, you can grant these third parties access to your AWS resources without sharing your AWS security credentials. Instead, the third party can access your AWS resources by assuming a role that you create in your AWS account.

Reference:

[http://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles\\_common-scenarios\\_third-party.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_common-scenarios_third-party.html)

**NEW QUESTION 32**

Your serverless architecture using AWS API Gateway, AWS Lambda, and AWS DynamoDB experienced a large increase in traffic to a sustained 400 requests per second, and dramatically increased in failure rates. Your requests, during normal operation, last 500 milliseconds on average. Your DynamoDB table did not exceed 50% of provisioned throughput, and Table primary keys are designed correctly. What is the most likely issue?

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

**Answer: C**

**Explanation:**

AWS API Gateway by default throttles at 500 requests per second steady-state, and 1000 requests per second at spike. Lambda, by default, throttles at 100 concurrent requests for safety. At 500 milliseconds (half of a second) per request, you can expect to support 200 requests per second at 100 concurrency. This is less than the 400 requests per second your system now requires. Make a limit increase request via the AWS Support Console.

AWS Lambda: Concurrent requests safety throttle per account -> 100

Reference: [http://docs.aws.amazon.com/general/latest/gr/aws\\_service\\_limits.html#limits\\_lambda](http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_lambda)

**NEW QUESTION 36**

For AWS CloudFormation, which stack state refuses UpdateStack calls?

- A. `UPDATE_ROLLBACK_FAILED`
- B. `UPDATE_ROLLBACK_COMPLETE`
- C. `UPDATE_COMPLETE`
- D. `CREATE_COMPLETE`

**Answer: A**

**Explanation:**

When a stack is in the UPDATE\_ROLLBACK\_FAILED state, you can continue rolling it back to return it to a working state (to UPDATE\_ROLLBACK\_COMPLETE). You cannot update a stack that is in the UPDATE\_ROLLBACK\_FAILED state. However, if you can continue to roll it back, you can return the stack to its original settings and try to update it again.

Reference:

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

**NEW QUESTION 39**

What is the scope of an EBS volume?

- A. VPC
- B. Region
- C. Placement Group
- D. Availability Zone

**Answer: D**

**Explanation:**

An Amazon EBS volume is tied to its Availability Zone and can be attached only to instances in the same Availability Zone.  
 Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/resources.html>

**NEW QUESTION 44**

You are experiencing performance issues writing to a DynamoDB table. Your system tracks high scores for video games on a marketplace. Your most popular game experiences all of the performance issues. What is the most likely problem?

- A. DynamoDB's vector clock is out of sync, because of the rapid growth in request for the most popular game.
- B. You selected the Game ID or equivalent identifier as the primary partition key for the table.
- C. Users of the most popular video game each perform more read and write requests than average.
- D. You did not provision enough read or write throughput to the tabl

**Answer: B**

**Explanation:**

The primary key selection dramatically affects performance consistency when reading or writing to DynamoDB. By selecting a key that is tied to the identity of the game, you forced DynamoDB to create a hotspot in the table partitions, and over-request against the primary key partition for the popular game. When it stores data, DynamoDB dMdes a table's items into multiple partitions, and distributes the data primarily based upon the partition key value. The provisioned throughput associated with a table is also dMded evenly among the partitions, with no sharing of provisioned throughput across partitions. Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GuidelinesForTables.html#Guideli nesForTables.UniformWorkload>

**NEW QUESTION 48**

You need to deploy a new application version to production. Because the deployment is high-risk, you need to roll the new version out to users over a number of hours, to make sure everything is working correctly. You need to be able to control the proportion of users seeing the new version of the application down to the percentage point.

You use ELB and EC2 with Auto Scaling Groups and custom AMIs with your code pre-installed assigned to Launch Configurations. There are no database-level changes during your deployment. You have been told you cannot spend too much money, so you must not increase the number of EC2 instances much at all during the deployment, but you also need to be able to switch back to the original version of code quickly if something goes wrong. What is the best way to meet these requirements?

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

**Answer: A**

**Explanation:**

Only Weighted Round Robin DNS Records and reverse proxies allow such fine-grained tuning of traffic splits. The Blue-Green option does not meet the requirement that we mitigate costs and keep overall EC2 fileet size consistent, so we must select the 2 ELB and ASG option with WRR DNS tuning. This method is called A/B deployment and/or Canary deployment.

Reference: <https://d0.awsstatic.com/whitepapers/overview-of-deployment-options-on-aws.pdf>

**NEW QUESTION 50**

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

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

**Answer: D**

**Explanation:**

You are not guaranteed 10gigabit performance, except within a placement group.

A placement group is a logical grouping of instances within a single Availability Zone. Using placement groups enables applications to participate in a low-latency, 10 Gbps network. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

**NEW QUESTION 55**

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

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

**Answer:** C

**Explanation:**

Use NotificationARNs.member.N when making a CreateStack call to push stack events into SNS in nearly real-time.

Reference:

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

**NEW QUESTION 60**

What does it mean if you have zero IOPS and a non-empty I/O queue for all EBS volumes attached to a running EC2 instance?

- A. The I/O queue is buffer flushing.
- B. Your EBS disk head(s) is/are seeking magnetic stripes.
- C. The EBS volume is unavailable.
- D. You need to re-mount the EBS volume in the OS.

**Answer:** C

**Explanation:**

This is the definition of Unavailable from the EC2 and EBS SLA.

"Unavailable" and "Unavailability" mean... For Amazon EBS, when all of your attached volumes perform zero read write IO, with pending IO in the queue.

Reference: <https://aws.amazon.com/ec2/sla/>

**NEW QUESTION 62**

From a compliance and security perspective, which of these statements is true?

- A. You do not ever need to rotate access keys for AWS IAM Users.
- B. You do not ever need to rotate access keys for AWS IAM Roles, nor AWS IAM Users.
- C. None of the other statements are true.
- D. You do not ever need to rotate access keys for AWS IAM Role

**Answer:** D

**Explanation:**

IAM Role Access Keys are auto-rotated by AWS on your behalf; you do not need to rotate them.

The application is granted the permissions for the actions and resources that you've defined for the role through the security credentials associated with the role.

These security credentials are temporary and we

rotate them automatically. We make new credentials available at least five minutes prior to the expiration of the old credentials.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

**NEW QUESTION 67**

Which of these is not a reason a Multi-AZ RDS instance will failover?

- A. An Availability Zone outage
- B. A manual failover of the DB instance was initiated using Reboot with failover
- C. To autoscale to a higher instance class
- D. The primary DB instance fails

**Answer:** C

**Explanation:**

The primary DB instance switches over automatically to the standby replica if any of the following conditions occur: An Availability Zone outage, the primary DB instance fails, the DB instance's server type is changed, the operating system of the DB instance is undergoing software patching, a manual failover of the DB instance was initiated using Reboot with failover

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

**NEW QUESTION 72**

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:**

The IAM Role must be fine, as if it were broken, NO jobs would be processed since the system would never be able to get any queue messages. The same reasoning applies to the routing table change. The scaling metric is fine, as instance count increased when the queue depth increased due to more messages entering than exiting. Thus, the only reasonable option is that some of the recent messages must be malformed and unprocessable.

Reference:

<https://github.com/andrew-templeton/cloudacademy/blob/fca920b45234bbe99cc0e8efb9c65134884dd489/questions/null>

**NEW QUESTION 74**

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

cost the most per month to operate?

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

**Answer: C**

**Explanation:**

Cost Allocation Tagging is a built-in feature of AWS, and when coupled with the Cost Explorer, provides a simple and robust way to track expenses. You can also use tags to filter views in Cost Explorer. Note that before you can filter views by tags in Cost Explorer, you must have applied tags to your resources and activate them, as described in the following sections. For more information about Cost Explorer, see Analyzing Your Costs with Cost Explorer. Reference: <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloc-tags.html>

**NEW QUESTION 79**

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

Which of these is the best system architectures for this system?

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

**Answer: B**

**Explanation:**

The Elastic Beanstalk option is incorrect because it requires a constantly-polling instance, which may break and costs money. The Lambda fleet option is incorrect because AWS Lambda does not support GPU usage. The OpsWorks stack option both requires a constantly-polling instance, and also requires complex timing and capacity planning logic. The CloudFormation option requires no polling, has no always-on instances, and allows arbitrarily fast processing by simply setting the instance count as high as needed. Reference: <http://docs.aws.amazon.com/lambda/latest/dg/current-supported-versions.html>

**NEW QUESTION 84**

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

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

**Answer: C**

**Explanation:**

You need to design for N+1 redundancy on Availability Zones.  $ZONE\_COUNT = (REQUIRED\_INSTANCES / INSTANCE\_COUNT\_PER\_ZONE) + 1$ . To minimize cost, spread the instances across as many possible zones as you can. By using a through e, you are allocating 5 zones. Using 2 instances, you have 10 total instances. If a single zone fails, you have 4 zones left, with 2 instances each, for a total of 8 instances. By spreading out as much as possible, you have increased cost by only 25% and significantly de-risked an availability zone failure.

Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html#concepts-regions-availability-zones>

**NEW QUESTION 85**

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 record.
- 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 it.
- F. Use the one without it when deploying to production.

**Answer:** A

**Explanation:**

The best way to do this is with one template, and a Condition on the resource. Route53 does not allow null strings for records.

Reference:

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

**NEW QUESTION 90**

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

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

**Answer:** D

**Explanation:**

Because the 100x read ratio is mostly driven by a small subset, with caching, only a roughly equal number of reads to writes will miss the cache, since the supermajority will hit the top 1% scores. Knowing we need to set the values roughly equal when using caching, we select AWS ElastiCache, because CloudFront cannot directly cache DynamoDB queries, and ElastiCache is an excellent in-memory cache for database queries, rather than a distributed proxy cache for content delivery.

One solution would be to cache these reads at the application layer. Caching is a technique that is used in many high-throughput applications, offloading read activity on hot items to the cache rather than to the database. Your application can cache the most popular items in memory, or use a product such as ElastiCache to do the same.

Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GuidelinesForTables.html#GuidelinesForTables.CachePopularItem>

**NEW QUESTION 91**

What is the scope of an EBS snapshot?

- A. Availability Zone
- B. Placement Group
- C. Region
- D. VPC

**Answer:** C

**Explanation:**

An EBS snapshot is tied to its region and can only be used to create volumes in the same region. You can copy a snapshot from one region to another. For more information, see Copying an Amazon EBS Snapshot.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/resources.html>

**NEW QUESTION 95**

Your CTO is very worried about the security of your AWS account. How best can you prevent hackers from completely hijacking your account?

- A. Use short but complex password on the root account and any administrators.
- B. Use AWS IAM Geo-Lock and disallow anyone from logging in except for in your city.
- C. Use MFA on all users and accounts, especially on the root account.
- D. Don't write down or remember the root account password after creating the AWS account.

**Answer:** C

**Explanation:**

For increased security, we recommend that you configure multi-factor authentication (MFA) to help protect your AWS resources. MFA adds extra security because it requires users to enter a unique authentication code from an approved authentication device or SMS text message when they access AWS websites or services.

Reference: [http://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_credentials\\_mfa.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_mfa.html)

**NEW QUESTION 96**

.....

## Thank You for Trying Our Product

### We offer two products:

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

2nd - Questions and Answers in PDF Format

### AWS-Certified-DevOps-Engineer-Professional Practice Exam Features:

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

**100% Actual & Verified — Instant Download, Please Click**  
[Order The AWS-Certified-DevOps-Engineer-Professional Practice Test Here](#)