

# Amazon AWS-DevOps-Engineer-Professional Exam

**Volume: 173 Questions**

Question: 1

What method should I use to author automation if I want to wait for a CloudFormation stack to finish completing in a script?

- A. Event subscription using SQS.
- B. Event subscription using SNS.
- C. Poll using `ListStacks` / `list-stacks`.
- D. Poll using `GetStackStatus` / `get-stack-status`.

Answer: C

Explanation:

Event driven systems are good for IFTTT logic, but only polling will make a script wait to complete. `ListStacks` / `list-stacks` is a real method, `GetStackStatus` / `get-stack-status` is not.

Reference: <http://docs.aws.amazon.com/cli/latest/reference/cloudformation/list-stacks.html>

Question: 2

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

## Amazon AWS-DevOps-Engineer-Professional Exam

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/>

Question: 3

You need to perform ad-hoc business analytics queries on well-structured data. Data comes in constantly at a high velocity. Your business intelligence team can understand SQL. What AWS service(s) should you look to first?

- A. Kinesis Firehose + RDS
- B. Kinesis Firehose + RedShift
- C. EMR using Hive
- D. EMR running Apache Spark

Answer: B

Explanation:

Kinesis Firehose provides a managed service for aggregating streaming data and inserting it into RedShift. RedShift also supports ad-hoc queries over well-structured data using a SQL-compliant wire protocol, so the business team should be able to adopt this system easily.

Reference: <https://aws.amazon.com/kinesis/firehose/details/>

Question: 4

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

## Amazon AWS-DevOps-Engineer-Professional Exam

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

Answer: B

Explanation:

Since access and storage for games is uniform, and you need to have ordering within each game for the scores (to access the highest value), your hash (partition) key should be the GameID, and there should be a range key for HighestScore.

Reference:

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

Question: 5

What is server immutability?

- A. Not updating a server after creation.
- B. The ability to change server counts.
- C. Updating a server after creation.
- D. The inability to change server counts.

Answer: A

Explanation:

... disposable upgrades offer a simpler way to know if your application has unknown dependencies. The underlying EC2 instance usage is considered temporary or ephemeral in nature for the period of deployment until the current release is active. During the new release, a new set of EC2 instances are rolled out by terminating older instances. This type of upgrade technique is more common in an immutable infrastructure.

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

Question: 6

You run a clustered NoSQL database on AWS EC2 using AWS EBS. You need to reduce latency

## Amazon AWS-DevOps-Engineer-Professional Exam

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>

Question: 7

Fill the blanks: \_\_\_\_\_ helps us track AWS API calls and transitions, \_\_\_\_\_ helps to understand what resources we have now, and \_\_\_\_\_ allows auditing credentials and logins.

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

Answer: C

Explanation:

You can use AWS CloudTrail to get a history of AWS API calls and related events for your account. This includes calls made by using the AWS Management Console, AWS SDKs, command line tools, and higher-level AWS services.

Reference:

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

## Amazon AWS-DevOps-Engineer-Professional Exam

Question: 8

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>

Question: 9

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:

The high-availability feature is not a scaling solution for read-only scenarios; you cannot use a standby replica to serve read traffic. To service read-only traffic, you should use a Read Replica. For more information, see Working with PostgreSQL, MySQL, and MariaDB Read Replicas.

Reference:

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

## Amazon AWS-DevOps-Engineer-Professional Exam

Question: 10

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>

Question: 11

You need your CI to build AMIs with code pre-installed on the images on every new code push. You need to do this as cheaply as possible. How do you do this?

- A. Bid on spot instances just above the asking price as soon as new commits come in, perform all instance configuration and setup, then create an AMI based on the spot instance.
- B. Have the CI launch a new on-demand EC2 instance when new commits come in, perform all instance configuration and setup, then create an AMI based on the on-demand instance.
- C. Purchase a Light Utilization Reserved Instance to save money on the continuous integration machine. Use these credits whenever you create AMIs on instances.
- D. When the CI instance receives commits, attach a new EBS volume to the CI machine. Perform all setup on this EBS volume so you don't need a new EC2 instance to create the AMI.

Answer: A

Explanation:

Spot instances are the cheapest option, and you can use minimum run duration if your AMI takes more than a few minutes to create.

## Amazon AWS-DevOps-Engineer-Professional Exam

Spot instances are also available to run for a predefined duration – in hourly increments up to six hours in length – at a significant discount (30-45%) compared to On-Demand pricing plus an additional 5% during off-peak times<sup>1</sup> for a total of up to 50% savings.

Reference: <https://aws.amazon.com/ec2/spot/pricing/>

Question: 12

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

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

Answer: A

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>

Question: 13

You need to process long-running jobs once and only once. How might you do this?

- A. Use an SNS queue and set the visibility timeout to long enough for jobs to process.
- B. Use an SQS queue and set the reprocessing timeout to long enough for jobs to process.
- C. Use an SQS queue and set the visibility timeout to long enough for jobs to process.
- D. Use an SNS queue and set the reprocessing timeout to long enough for jobs to process.

Answer: C

Explanation:

The message timeout defines how long after a successful receive request SQS waits before

## Amazon AWS-DevOps-Engineer-Professional Exam

allowing jobs to be seen by other components, and proper configuration prevents duplicate processing.

Reference:

<http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/MessageLifecycle.html>

Question: 14

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 tables. Lambda will scale rapidly enough for the traffic spikes.

B. Use a CloudFront distribution with access log delivery to S3. Clicks should be recorded as querystring GETs to the distribution. Reports are built and sent by periodically running EMR jobs over the access logs in S3.

C. Use API Gateway invoking Lambdas which PutRecords into Kinesis, and EMR running Spark performing GetRecords on Kinesis to scale with spikes. Spark on EMR outputs the analysis to S3, which are sent out via email.

D. Use AWS Elasticsearch service and EC2 Auto Scaling groups. The Autoscaling groups scale based on click throughput and stream into the Elasticsearch domain, which is also scalable. 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/>

Question: 15

Your system automatically provisions EIPs to EC2 instances in a VPC on boot. The system provisions the whole VPC and stack at once. You have two of them per VPC. On your new AWS

## Amazon AWS-DevOps-Engineer-Professional Exam

account, your attempt to create a Development environment failed, after successfully creating Staging and Production environments in the same region. What happened?

- A. You didn't choose the Development version of the AMI you are using.
- B. You didn't set the Development flag to true when deploying EC2 instances.
- C. You hit the soft limit of 5 EIPs per region and requested a 6th.
- D. You hit the soft limit of 2 VPCs per region and requested a 3rd.

Answer: C

Explanation:

There is a soft limit of 5 EIPs per Region for VPC on new accounts. The third environment could not allocate the 6th EIP.

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

Question: 16

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/>

Question: 17

How does Amazon RDS multi Availability Zone model work?

## Amazon AWS-DevOps-Engineer-Professional Exam

- 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>

Question: 18

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>

Question: 19