



## Amazon

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

Amazon AWS Certified DevOps Engineer Professional

## About ExamBible

### *Your Partner of IT Exam*

## Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

### NEW 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>

### NEW 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 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 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/>

### NEW QUESTION 4

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/awscloudtrail/latest/userguide/cloudtrail-user-guide.html>

### NEW QUESTION 5

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
- D. Use these credits whenever you create AMIs on instances.
- E. When the CI instance receives commits, attach a new EBS volume to the CI machine
- F. 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.

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 for a total of up to 50% savings.

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

**NEW QUESTION 6**

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 spiky, 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 distribution
- 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 scalable
- 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 7**

What is the scope of an EC2 security group?

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

**Answer:** C

**Explanation:**

A security group is tied to a region and can be assigned only to instances in the same region. You can't enable an instance to communicate with an instance outside its region using security group rules. Traffic from an instance in another region is seen as WAN bandwidth.

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

**NEW QUESTION 8**

You run accounting software in the AWS cloud. This software needs to be online continuously during the day every day of the week, and has a very static requirement for compute resources. You also have other, unrelated batch jobs that need to run once per day at any time of your choosing. How should you minimize cost?

- A. Purchase a Heavy Utilization Reserved Instance to run the accounting software
- B. Turn it off after hour
- C. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.
- D. Purchase a Medium Utilization Reserved Instance to run the accounting software
- E. Turn it off after hour
- F. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.
- G. Purchase a Light Utilization Reserved Instance to run the accounting software
- H. Turn it off after hour
- I. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.
- J. Purchase a Full Utilization Reserved Instance to run the accounting software
- K. Turn it off after hour
- L. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.

**Answer:** A

**Explanation:**

Because the instance will always be online during the day, in a predictable manner, and there are a sequence of batch jobs to perform at any time, we should run the batch jobs when the accounting software is off. We can achieve Heavy Utilization by alternating these times, so we should purchase the reservation as such, as this represents the lowest cost. There is no such thing as a "Full" level utilization purchase on EC2.

Reference: [https://d0.awsstatic.com/whitepapers/Cost\\_Optimization\\_with\\_AWS.pdf](https://d0.awsstatic.com/whitepapers/Cost_Optimization_with_AWS.pdf)

#### NEW QUESTION 9

Which EBS volume type is best for high performance NoSQL cluster deployments?

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

**Answer:** A

#### Explanation:

io1 volumes, or Provisioned IOPS (PIOPS) SSDs, are best for: Critical business applications that require sustained IOPS performance, or more than 10,000 IOPS or 160 MiB/s of throughput per volume, like large database workloads, such as MongoDB.

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

#### NEW QUESTION 10

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 10

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 15

You want to pass queue messages that are 1GB each. How should you achieve this?

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

**Answer:** B

#### Explanation:

You can manage Amazon SQS messages with Amazon S3. This is especially useful for storing and retrieving messages with a message size of up to 2 GB. To manage Amazon SQS messages with Amazon S3, use the Amazon SQS Extended Client Library for Java.

Reference:

<http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/s3-messages.html>

#### NEW QUESTION 16

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 20**

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 21**

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 22**

Which deployment method, when using AWS Auto Scaling Groups and Auto Scaling Launch Configurations, enables the shortest time to live for individual servers?

- A. Pre-baking AMIs with all code and configuration on deploys.
- B. Using a Dockerfile bootstrap on instance launch.
- C. Using UserData bootstrapping scripts.
- D. Using AWS EC2 Run Commands to dynamically SSH into fileset

**Answer:** A

**Explanation:**

Note that the bootstrapping process can be slower if you have a complex application or multiple applications to install. Managing a fileset 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. Prebaking is a process of embedding a significant portion of your application artifacts within your base AMI. During the deployment process you can customize application installations by using EC2 instance artifacts such as instance tags, instance metadata, and Auto Scaling groups.

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

**NEW QUESTION 26**

Which major database needs a BYO license?

- A. PostgreSQL
- B. MariaDB
- C. MySQL
- D. Oracle

**Answer:** D

**Explanation:**

Oracle is not open source, and requires a bring your own license model.

Reference: [http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_Oracle.html](http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Oracle.html)

**NEW QUESTION 27**

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 32**

You need to know when you spend \$1000 or more on AWS. What's the easy way for you to see that notification?

- A. AWS CloudWatch Events tied to API calls, when certain thresholds are exceeded, publish to SNS.
- B. Scrape the billing page periodically and pump into Kinesis.
- C. AWS CloudWatch Metrics + Billing Alarm + Lambda event subscription
- D. When a threshold is exceeded, email the manager.
- E. Scrape the billing page periodically and publish to SNS

**Answer:** C

**Explanation:**

Even if you're careful to stay within the free tier, it's a good idea to create a billing alarm to notify you if you exceed the limits of the free tier. Billing alarms can help to protect you against unknowingly accruing charges if you inadvertently use a service outside of the free tier or if traffic exceeds your expectations. Reference: <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/free-tier-alarms.html>

**NEW QUESTION 36**

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 37**

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

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

**Answer:** A

**Explanation:**

You must use CloudTrail Log File Validation (default or custom implementation), as any other tracking method is subject to forgery in the event of a full account compromise by sophisticated enough hackers. Validated log files are invaluable in security and forensic investigations. For example, a validated log file enables you to assert positively that the log file itself has not changed, or that particular user credentials performed specific API activity. The CloudTrail log file integrity validation process also lets you know if a log file has been deleted or changed, or assert positively that no log files were delivered to your account during a given period of time.

Reference:

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

**NEW QUESTION 39**

Which of these is not a Pseudo Parameter in AWS CloudFormation?

- A. AWS::StackName
- B. AWS::AccountId
- C. AWS::StackArn
- D. AWS::NotificationARNs

**Answer:** C

**Explanation:**

This is the complete list of Pseudo Parameters: AWS::AccountId, AWS::NotificationARNs, AWS::NoValue, AWS::Region, AWS::StackId, AWS::StackName

Reference:

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

**NEW QUESTION 41**

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 table.

**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 divides 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 divided evenly among the partitions, with no sharing of provisioned throughput across partitions. Reference:

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

**NEW QUESTION 47**

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

- A. Fn::Split
- B. Fn::FindInMap
- C. Fn::Select
- D. Fn::GetAZs

**Answer: A**

**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 51**

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 API
- 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 regions
- G. Make sure both regions use Auto Scaling Groups behind ELBs.

**Answer: D**

**Explanation:**

Standard volumes, or Magnetic volumes, are best for: Cold workloads where data is infrequently accessed, or scenarios where the lowest storage cost is important.

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

**NEW QUESTION 56**

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

- A. Use a CloudFormation Custom Resource Template by selecting an API call to proxy for create, update, and delete action

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

**Answer:** D

**Explanation:**

Custom resources provide a way for you to write custom provisioning logic in AWS CloudFormation template and have AWS CloudFormation run it during a stack operation, such as when you create, update or delete a stack. For more information, see Custom Resources.

Reference:

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

**NEW QUESTION 59**

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

- A. Use a large AWS Directory Service Simple AD.
- B. Use a large AWS Directory Service AD Connector.
- C. Use an Sync Domain running on AWS Directory Service.
- D. Use an AWS Directory Sync Domain running on AWS Lambda

**Answer:** B

**Explanation:**

You must use AD Connector as a power-user of Microsoft Active Directory. Simple AD only works with a subset of AD functionality. Sync Domains do not exist; they are made up answers.

AD Connector is a directory gateway that allows you to proxy directory requests to your on-premises Microsoft Active Directory, without caching any information in the cloud. AD Connector comes in 2 sizes; small and large. A small AD Connector is designed for smaller organizations of up to 500 users. A large AD Connector is designed for larger organizations of up to 5,000 users.

Reference: <https://aws.amazon.com/directoryservice/details/>

**NEW QUESTION 60**

Which of these is not a CloudFormation Helper Script?

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

**Answer:** C

**Explanation:**

This is the complete list of CloudFormation Helper Scripts: cfn-init, cfn-signal, cfn-get-metadata, cfn-hup Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-helper-scripts-reference.html>

**NEW QUESTION 61**

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

- A. AWS CloudFormation
- B. AWS OpsWorks
- C. AWS ELB + EC2 with CLI Push
- D. AWS Elastic Beanstalk

**Answer:** D

**Explanation:**

Elastic Beanstalk's primary mode of operation exactly supports this use case out of the box. It is simpler than all the other options for this question.

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.

Reference: [http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create\\_deploy\\_Ruby\\_rails.html](http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_Ruby_rails.html)

**NEW QUESTION 66**

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 without needing to worry about scaling expensive uploads processes, authentication/authorization and so forth?

- 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 SANIL compliant systems to build authorization policie
- D. Users log in with a username and password, and are given a token they can use indefinitely to make calls against the photo infrastructure.
- E. Use AWS API Gateway with a constantly rotating API Key to allow access from the client-sid

- 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:**

The short answer is that Amazon Cognito is a superset of the functionality provided by web identity federation. It supports the same providers, and you configure your app and authenticate with those providers in the same way. But Amazon Cognito includes a variety of additional features. For example, it enables your users to start using the app as a guest user and later sign in using one of the supported identity providers.

Reference:

<https://blogs.aws.amazon.com/security/post/Tx3SYCORF5EKRCO/How-Does-Amazon-Cognito-Relate-to-Existing-Web-Identity-Federation/>

**NEW QUESTION 70**

Your CTO has asked you to make sure that you know what all users of your AWS account are doing to change resources at all times. She wants a report of who is doing what over time, reported to her once per week, for as broad a resource type group as possible. How should you do this?

- A. Create a global AWS CloudTrail Trail
- B. Configure a script to aggregate the log data delivered to S3 once per week and deliver this to the CTO.
- C. Use CloudWatch Events Rules with an SNS topic subscribed to all AWS API call
- D. Subscribe the CTO to an email type delivery on this SNS Topic.
- E. Use AWS IAM credential reports to deliver a CSV of all uses of IAM User Tokens over time to the CTO.
- F. Use AWS Config with an SNS subscription on a Lambda, and insert these changes over time into a DynamoDB table
- G. Generate reports based on the contents of this table.

**Answer:** A

**Explanation:**

This is the ideal use case for AWS CloudTrail.

CloudTrail provides visibility into user activity by recording API calls made on your account. CloudTrail records important information about each API call, including the name of the API, the identity of the caller, the time of the API call, the request parameters, and the response elements returned by the AWS service. This information helps you to track changes made to your AWS resources and to troubleshoot operational issues. CloudTrail makes it easier to ensure compliance with internal policies and regulatory standards. Reference: <https://aws.amazon.com/Cloudtrail/faqs/>

**NEW QUESTION 72**

Which is not a restriction on AWS EBS Snapshots?

- A. Snapshots which are shared cannot be used as a basis for other snapshots.
- B. You cannot share a snapshot containing an AWS Access Key ID or AWS Secret Access Key.
- C. You cannot share unencrypted snapshots.
- D. Snapshot restorations are restricted to the region in which the snapshots are created

**Answer:** A

**Explanation:**

Snapshots shared with other users are usable in full by the recipient, including but limited to the ability to base modified volumes and snapshots.

Reference:

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

**NEW QUESTION 76**

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

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

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

**Answer:** A

**Explanation:**

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 fleet 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 78

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:

All suggested periodic options are sensitive to server failure during or between periodic flushes. Streaming to Lambda and then logging to CloudWatch Logs will make the system resilient to instance and Availability Zone failures.

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

#### NEW QUESTION 79

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

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

**Answer:** B

#### Explanation:

S3 stores all snapshots. If S3 is unavailable, snapshots are unavailable.

Amazon EC2 also uses Amazon S3 to store snapshots (backup copies) of the data volumes. You can use snapshots for recovering data quickly and reliably in case of application or system failures. You can also use snapshots as a baseline to create multiple new data volumes, expand the size of an existing data volume, or move data volumes across multiple Availability Zones, thereby making your data usage highly scalable. For more information about using data volumes and snapshots, see Amazon Elastic Block Store.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonS3.html>

#### NEW QUESTION 82

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

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

**Answer:** B

#### Explanation:

To declaratively model build and destroy of a cluster, you need to use AWS CloudFormation. OpsWorks and Elastic Beanstalk cannot directly model EMR Clusters. The CLI is not declarative, and CLI Composer does not exist.

Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-emr-cluster.html>

#### NEW QUESTION 87

What is true of the way that encryption works with EBS?

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

**Answer:** C

#### Explanation:

Snapshots that are taken from encrypted volumes are automatically encrypted. Volumes that are created from encrypted snapshots are also automatically encrypted. Your encrypted volumes and any associated snapshots always remain protected. For more information, see Amazon EBS Encryption.

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

#### NEW QUESTION 88

When thinking of AWS OpsWorks, which of the following is true?

- A. Stacks have many layers, layers have many instances.
- B. Instances have many stacks, stacks have many layers.
- C. Layers have many stacks, stacks have many instances.
- D. Layers have many instances, instances have many stack

**Answer:** A

**Explanation:**

The stack is the core AWS OpsWorks component. It is basically a container for AWS resources—Amazon EC2 instances, Amazon RDS database instances, and so on—that have a common purpose and should be logically managed together. You define the stack's constituents by adding one or more layers. A layer represents a set of Amazon EC2 instances that serve a particular purpose, such as serving applications or hosting a database server. An instance represents a single computing resource, such as an Amazon EC2 instance.

Reference: <http://docs.aws.amazon.com/opsworks/latest/userguide/welcome.html>

**NEW QUESTION 89**

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 laye
- D. When the instances are booted and the HTTP APIs are up, submit processing requests to indMdual instances.
- E. Make an S3 notification configuration which publishes to AWS Lambda on the manifest bucke
- 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 instanc
- G. Create an SQS queue of the images in the manifes
- 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++ extensio
- J. Make an S3 notification configuration on the manifest, which publishes to another AWS Lambda running controller cod
- K. This controller code publishes all the images in the manifest to AWS Kinesi
- L. Your ANN code Lambda Function uses the Kinesis as an Event Sourc
- M. The system automatically scales when the stream contains image events.
- N. Create an Auto Scaling, Load Balanced Elastic Beanstalk worker tier Application and Environmen
- O. Deploy the ANN code to G2 instances in this tie
- 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 fileet 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 92**

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, you 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, inclusiv

**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 though 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 96**

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 `<code>Parameter</code>` for `<code>environment</code>`, and add a `<code>Condition</code>` on the Route53 `<code>Resource</code>` 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 99**

What is web identity federation?

- A. Use of an identity provider like Google or Facebook to become an AWS IAM User.
- B. Use of an identity provider like Google or Facebook to exchange for temporary AWS security credentials.
- C. Use of AWS IAM User tokens to log in as a Google or Facebook user.
- D. Use of AWS STS Tokens to log in as a Google or Facebook user.

**Answer:** B

**Explanation:**

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.

Reference: [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 104**

When thinking of AWS Elastic Beanstalk, which statement is true?

- A. Worker tiers pull jobs from SNS.
- B. Worker tiers pull jobs from HTTP.
- C. Worker tiers pull jobs from JSON.
- D. Worker tiers pull jobs from SQS.

**Answer:** D

**Explanation:**

Elastic Beanstalk installs a daemon on each Amazon EC2 instance in the Auto Scaling group to process Amazon SQS messages in the worker environment. The daemon pulls data off the Amazon SQS queue, inserts it into the message body of an HTTP POST request, and sends it to a user-configurable URL path on the local host. The content type for the message body within an HTTP POST request is `application/json` by default.

Reference:

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.html>

**NEW QUESTION 106**

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:**

Systems will always over-scale unless you choose the metric that runs out first and becomes constrained first. You also need to set the thresholds of the metric based on whether or not latency is affected by the change, to justify adding capacity instead of wasting money.

Reference: [http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/policy\\_creating.html](http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/policy_creating.html)

**NEW QUESTION 110**

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. Configure Rolling Deployments.
- B. Configure Enhanced Health Reporting.
- C. Configure Blue-Green Deployments.
- D. Configure a Dead Letter Queue.

**Answer:** D

**Explanation:**

Elastic Beanstalk worker environments support Amazon Simple Queue Service (SQS) dead letter queues. A dead letter queue is a queue where other (source)

queues can send messages that for some reason could not be successfully processed. A primary benefit of using a dead letter queue is the ability to sideline and isolate the unsuccessfully processed messages. You can then analyze any messages sent to the dead letter queue to try to determine why they were not successfully processed. Reference: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.html#worker-deadletter>

#### NEW QUESTION 115

You have a high security requirement for your AWS accounts. What is the most rapid and sophisticated setup you can use to react to AWS API calls to your account?

- A. Subscription to AWS Config via an SNS Topic
- B. Use a Lambda Function to perform in-flight analysis and react to changes as they occur.
- C. Global AWS CloudTrail setup delivering to S3 with an SNS subscription to the deliver notifications, pushing into a Lambda, which inserts records into an ELK stack for analysis.
- D. Use a CloudWatch Rule ScheduleExpression to periodically analyze IAM credential log
- E. Push the deltas for events into an ELK stack and perform ad-hoc analysis there.
- F. CloudWatch Events Rules which trigger based on all AWS API calls, submitting all events to an AWS Kinesis Stream for arbitrary downstream analysis.

**Answer:** D

#### **Explanation:**

CloudWatch Events allow subscription to AWS API calls, and direction of these events into Kinesis Streams. This allows a unified, near real-time stream for all API calls, which can be analyzed with any tool(s) of your choosing downstream.

Reference: [http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/EventTypes.html#api\\_event\\_type](http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/EventTypes.html#api_event_type)

#### NEW QUESTION 118

.....

## Relate Links

**100% Pass Your AWS-Certified-DevOps-Engineer-Professional Exam with Exambible Prep Materials**

<https://www.exambible.com/AWS-Certified-DevOps-Engineer-Professional-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>