AWS-Certified-DevOps-Engineer-Professional Dumps

Amazon AWS Certified DevOps Engineer Professional

https://www.certleader.com/AWS-Certified-DevOps-Engineer-Professional-dumps.html
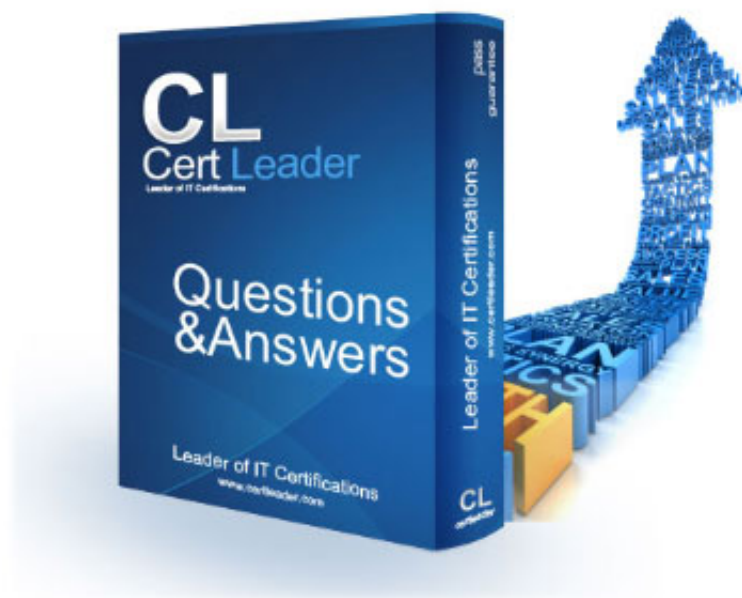
NEW QUESTION 1
What is server immutability?

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 count

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

NEW QUESTION 2
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 machin
D. Use these credits whenever your create AMIs on instances.
E. When the CI instance receives commits, attach a new EBS volume to the CI machin
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 timesl for a total of up to 50% savings.
Reference: https://aws.amazon.com/ec2/spot/pricing/

NEW QUESTION 3
Which of these is not an instrinsic 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::Se|ect, Ref
Reference:
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference.html

NEW QUESTION 4
Which status represents a failure state in AWS CIoudFormation?

A. <code>UPDATE_COMPLETE_CLEANUP_IN_PROGRESS</code>
B. <code>DELETE_COMPLETE_WITH_ARTIFACTS</code>
C. <code>ROLLBACK_IN_PROGRESS</code>
D. <code>ROLLBACK_FAILED</code>

Answer: C

Explanation:
ROLLBACK_IN_PROGRESS means an UpdateStack operation failed and the stack is in the process of trying to return to the valid, pre-update state.
UPDATE_COMPLETE_CLEANUP_IN_PROGRESS means an update was successful, and CIoudFormation is deleting any replaced, no longer used resources.
ROLLBACK_FA|LED is not a CloudFormation state (but UPDATE_ROLLBACK_FAILED is). DELETE_COMPLETE_W|TH_ART|FACTS does not exist at all.
Reference:
http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/using-cfn-updating-stacks.html

NEW QUESTION 5
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.htmI

NEW QUESTION 6
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/DeveIoperGuide/AutoScaIingGroupLifecycIe.html

NEW QUESTION 7
You want to pass queue messages that are 1GB each. How should you achieve this?

A. Use Kinesis as a buffer stream for message bodie
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 mediu
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/AWSSimpIeQueueService/latest/SQSDeveIoperGuide/s3-messages.html

NEW QUESTION 8
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 tabl

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.htmI

NEW QUESTION 9
Which major database needs a BYO license?

A. PostgreSQL
B. NIariaDB
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_OracIe.htmI

NEW QUESTION 10
What is the maximum supported single-volume throughput on EBS?

A. 320IV|iB/s
B. 160MiB/s
C. 40MiB/s
D. 640MiB/s

Answer: A

Explanation:
The ceiling throughput for PIOPS on EBS is 320MiB/s.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVo|umeTypes.htm| IIIIIEZIIII HWS-IIEIIII|]S-EII§iII|}|}I'-PI'0I|}SSi0IIilI EIIEIII

NEW QUESTION 10
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 bucke
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 accoun
E. Grant the associated IAM role full access to the bucke
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 PUT 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

NEW QUESTION 12
You meet once per month with your operations team to review the past month's data. During the meeting, you realize that 3 weeks ago, your monitoring system which pings over HTTP from outside AWS recorded a large spike in latency on your 3-tier web service API.
You use DynamoDB for the database layer, ELB, EBS, and EC2 for the business logic tier, and SQS, ELB, and EC2 for the presentation layer.
Which of the following techniques will NOT help you figure out what happened?

A. Check your CloudTrail log history around the spike's time for any API calls that caused slowness.
B. Review CloudWatch Metrics graphs to determine which component(s) slowed the system down.
C. Review your ELB access logs in S3 to see if any ELBs in your system saw the latency.
D. Analyze your logs to detect bursts in traffic at that tim

Answer: B

Explanation:
Metrics data are available for 2 weeks. If you want to store metrics data beyond that duration, you can retrieve it using our GetMetricStatistics API as well as a number of applications and tools offered by AWS partners.
Reference: https://aws.amazon.com/cloudwatch/faqs/

NEW QUESTION 13
You need to replicate API calls across two systems in real time. What tool should you use as a buffer and transport mechanism for API call events?

A. AWS SQS
B. AWS Lambda
C. AWS Kinesis
D. AWS SNS

Answer: C

Explanation:
AWS Kinesis is an event stream service. Streams can act as buffers and transport across systems for in-order programmatic events, making it ideal for replicating API calls across systems.
A typical Amazon Kinesis Streams application reads data from an Amazon Kinesis stream as data records. These applications can use the Amazon Kinesis Client Library, and they can run on Amazon EC2 instances. The processed records can be sent to dashboards, used to generate alerts, dynamically change pricing and advertising strategies, or send data to a variety of other AWS services. For information about Streams features and pricing, see Amazon Kinesis Streams.
Reference: http://docs.aws.amazon.com/kinesis/latest/dev/introduction.html

NEW QUESTION 18
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 create

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 21
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 O

Answer: C

Explanation:
This is the definition of Unavailable from the EC2 and EBS SLA.
"Unavailable" and "Unavai|abi|ity" 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/s|a/

NEW QUESTION 26
You need to create an audit log of all changes to customer banking data. You use DynamoDB to store this customer banking data. |t'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 Lambd
B. Log the changes toAWS CloudWatch Logs, removing sensitive information before logging.
C. Before writing to DynamoDB, do a pre-write acknoledgment to disk on the application sewer, removing sensitive information before loggin
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 acknoledgment to disk on the application sewer, removing sensitive information before loggin
H. Periodically pipe these files into CloudWatch Logs.

Answer: A

Explanation:
All suggested periodic options are sensitive to sewer 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 28
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 newjobs 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 correctl

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/fca920b45234bbe99cc0e8efb9c65134884dd48 9/questions/null

NEW QUESTION 33
Which of the following tools does not directly support AWS OpsWorks, for monitoring your stacks?

A. AWS Config
B. Amazon CloudWatch Nletrics
C. AWS CloudTrail
D. Amazon CloudWatch Logs

Answer: A

Explanation:
You can monitor your stacks in the following ways: AWS OpsWorks uses Amazon CloudWatch to provide thirteen custom metrics with detailed monitoring for each instance in the stack; AWS OpsWorks integrates with AWS CloudTrail to log every AWS OpsWorks API call and store the data in an Amazon S3 bucket; You can use Amazon CloudWatch Logs to monitor your stack's system, application, and custom logs. Reference:
http://docs.aws.amazon.com/opsworks/latest/userguide/monitoring.html

NEW QUESTION 38
What is a circular dependency in AWS CloudFormation?

A. When a Template references an earlier version of itself.
B. When Nested Stacks depend on each other.

C. When Resources form a DependOn loop.
D. When a Template references a region, which references the original Templat

Answer: C

Explanation:
To resolve a dependency error, add a DependsOn attribute to resources that depend on other resources in your template. In some cases, you must explicitly declare dependencies so that AWS CloudFormation can create or delete resources in the correct order. For example, if you create an Elastic IP and a VPC with an Internet gateway in the same stack, the Elastic IP must depend on the Internet gateway attachment. For additional information, see DependsOn Attribute.
Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/troubleshooting.htm|#troub|eshootin g-errors-dependence-error

NEW QUESTION 43
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 48
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.htmI

NEW QUESTION 49
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/weIcome.html

NEW QUESTION 51
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 55
Your company needs to automate 3 layers of a large cloud deployment. You want to be able to track this deployment's evolution as it changes over time, and carefully control any alterations. What is a good way to automate a stack to meet these requirements?

A. Use OpsWorks Stacks with three layers to model the layering in your stack.
B. Use CloudFormation Nested Stack Templates, with three child stacks to represent the three logicallayers of your cloud.
C. Use AWS Config to declare a configuration set that AWS should roll out to your cloud.
D. Use Elastic Beanstalk Linked Applications, passing the important DNS entires between layers using the metadata interface.

Answer: B

Explanation:
Only CloudFormation allows source controlled, declarative templates as the basis for stack automation. Nested Stacks help achieve clean separation of layers while simultaneously providing a method to control all layers at once when needed.
Reference:
https://blogs.aws.amazon.com/application-management/post/TxIT9JYOOS8AB9I/Use-Nested-Stacks-to- Create-Reusable-Templates-and-Support-Role-Specialization

NEW QUESTION 59
......

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

   All transactions are protected by VeriSign!

100% Pass Your AWS-Certified-DevOps-Engineer-Professional Exam with Our Prep Materials Via below:

https://www.certleader.com/AWS-Certified-DevOps-Engineer-Professional-dumps.html