



BDS-C00^{Q&As}

AWS Certified Big Data - Speciality (BDS-C00)

Pass Amazon BDS-C00 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.lead4pass.com/aws-certified-big-data-specialty.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

How should an Administrator BEST architect a large multi-layer Long Short-Term Memory (LSTM) recurrent neural network (RNN) running with MXNET on Amazon EC2? (Choose two.)

- A. Use data parallelism to partition the workload over multiple devices and balance the workload within the GPUs.
- B. Use compute-optimized EC2 instances with an attached elastic GPU.
- C. Use general purpose GPU computing instances such as G3 and P3.
- D. Use processing parallelism to partition the workload over multiple storage devices and balance the workload within the GPUs.

Correct Answer: AC

QUESTION 2

A data engineer is running a DWH on a 25-node Redshift cluster of a SaaS service. The data engineer needs to build a dashboard that will be used by customers. Five big customers represent 80% of usage, and there is a long tail of dozens of smaller customers. The data engineer has selected the dashboarding tool.

How should the data engineer make sure that the larger customer workloads do NOT interfere with the smaller customer workloads?

- A. Apply query filters based on customer-id that can NOT be changed by the user and apply distribution keys on customer-id.
- B. Place the largest customers into a single user group with a dedicated query queue and place the rest of the customers into a different query queue.
- C. Push aggregations into an RDS for Aurora instance. Connect the dashboard application to Aurora rather than Redshift for faster queries.
- D. Route the largest customers to a dedicated Redshift cluster. Raise the concurrency of the multi-tenant Redshift cluster to accommodate the remaining customers.

Correct Answer: D

QUESTION 3

A user has deployed an application on his private cloud. The user is using his own monitoring tool. He wants to configure that whenever there is an error, the monitoring tool should notify him via SMS.

Which of the below mentioned AWS services will help in this scenario?

- A. None because the user infrastructure is in the private cloud/
- B. AWS SNS
- C. AWS SES



D. AWS SMS

Correct Answer: B

QUESTION 4

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

Correct Answer: A

QUESTION 5

An advertising organization uses an application to process a stream of events that are received from clients in multiple unstructured formats.

The application does the following:

Transforms the events into a single structured format and streams them to Amazon Kinesis for real-time analysis.

Stores the unstructured raw events from the log files on local hard drivers that are rotated and uploaded to Amazon S3.

The organization wants to extract campaign performance reporting using an existing Amazon redshift cluster.

Which solution will provide the performance data with the LEAST number of operations?

- A. Install the Amazon Kinesis Data Firehose agent on the application servers and use it to stream the log files directly to Amazon Redshift.
- B. Create an external table in Amazon Redshift and point it to the S3 bucket where the unstructured raw events are stored.
- C. Write an AWS Lambda function that triggers every hour to load the new log files already in S3 to Amazon redshift.
- D. Connect Amazon Kinesis Data Firehose to the existing Amazon Kinesis stream and use it to stream the event directly to Amazon Redshift.

Correct Answer: B



QUESTION 6

Does AWS Direct Connect allow you access to all Availabilities Zones within a Region?

- A. Depends on the type of connection
- B. No
- C. Yes
- D. Only when there's just one availability zone in a region. If there are more than one, only one availability zone can be accessed directly.

Correct Answer: A

QUESTION 7

A data engineer wants to use an Amazon Elastic Map Reduce for an application. The data engineer needs to make sure it complies with regulatory requirements. The auditor must be able to confirm at any point which servers are running and

which network access controls are deployed.

Which action should the data engineer take to meet this requirement?

- A. Provide the auditor IAM accounts with the SecurityAudit policy attached to their group.
- B. Provide the auditor with SSH keys for access to the Amazon EMR cluster.
- C. Provide the auditor with CloudFormation templates.
- D. Provide the auditor with access to AWS DirectConnect to use their existing tools.

Correct Answer: C

QUESTION 8

You are configuring your company's application to use Auto Scaling and need to move user state information. Which of the following AWS services provides a shared data store with durability and low latency?

- A. Amazon Simple Storage Service
- B. Amazon DynamoDB
- C. Amazon EC2 instance storage
- D. AWS ElasticCache Memcached

Correct Answer: A



QUESTION 9

A data engineer is about to perform a major upgrade to the DDL contained within an Amazon Redshift cluster to support a new data warehouse application. The upgrade scripts will include user permission updates, view and table structure changes as well as additional loading and data manipulation tasks.

The data engineer must be able to restore the database to its existing state in the event of issues.

Which action should be taken prior to performing this upgrade task?

- A. Run an UNLOAD command for all data in the warehouse and save it to S3.
- B. Create a manual snapshot of the Amazon Redshift cluster.
- C. Make a copy of the automated snapshot on the Amazon Redshift cluster.
- D. Call the `waitForSnapshotAvailable` command from either the AWS CLI or an AWS SDK.

Correct Answer: B

Reference: <https://docs.aws.amazon.com/redshift/latest/mgmt/working-with-snapshots.html#working-with-snapshot-restore-table-from-snapshot>

QUESTION 10

By default what are ENIs that are automatically created and attached to instances using the EC2 console set to do when the attached instance terminates?

- A. Remain as is
- B. Terminate
- C. Hibernate
- D. Pause

Correct Answer: B

QUESTION 11

You currently run your infrastructure on Amazon EC2 instances behind on Auto Scaling group. All logs for your application are currently written to ephemeral storage. Recently your company experienced a major bug in code that made it through testing and was ultimately deployed to your fleet. This bug triggered your Auto Scaling group to scale up and back down before you could successfully retrieve the logs off your server to better assist you in troubleshooting the bug.

Which technique should you use to make sure you are able to review your logs after your instances have shut down?

- A. Configure the ephemeral policies on your Auto Scaling group to back up on terminate



- B. Configure your Auto Scaling policies to create a snapshot of all ephemeral storage on terminate
- C. Install the CloudWatch logs Agent on your AMI, and configure CloudWatch Logs Agent to stream your logs
- D. Install the CloudWatch monitoring agent on your AMI, and set up a new SNS alert for CloudWatch metrics that triggers the CloudWatch monitoring agent to backup all logs on the ephemeral drive
- E. Install the CloudWatch Logs Agent on your AMI. Update your Scaling policy to enable automated CloudWatch Log copy

Correct Answer: C

QUESTION 12

An Amazon EMR cluster using EMRFS has access to petabytes of data on Amazon S3, originating from multiple unique data sources. The customer needs to query common fields across some of the data sets to be able to perform interactive joins and then display results quickly.

Which technology is most appropriate to enable this capability?

- A. Presto
- B. MicroStrategy
- C. Pig
- D. R Studio

Correct Answer: C

QUESTION 13

An organization uses a custom map reduce application to build monthly reports based on many small data files in an Amazon S3 bucket. The data is submitted from various business units on a frequent but unpredictable schedule. As the

dataset continues to grow, it becomes increasingly difficult to process all of the data in one day. The organization has scaled up its Amazon EMR cluster, but other optimizations could improve performance.

The organization needs to improve performance with minimal changes to existing processes and applications.

What action should the organization take?

- A. Use Amazon S3 Event Notifications and AWS Lambda to create a quick search file index in DynamoDB.
- B. Add Spark to the Amazon EMR cluster and utilize Resilient Distributed Datasets in-memory.
- C. Use Amazon S3 Event Notifications and AWS Lambda to index each file into an Amazon Elasticsearch Service cluster.
- D. Schedule a daily AWS Data Pipeline process that aggregates content into larger files using S3DistCp.
- E. Have business units submit data via Amazon Kinesis Firehose to aggregate data hourly into Amazon S3.



Correct Answer: B

[BDS-C00 VCE Dumps](#)

[BDS-C00 Study Guide](#)

[BDS-C00 Braindumps](#)



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !

- 100% Guaranteed Success
- 100% Money Back Guarantee
- 365 Days Free Update
- Instant Download After Purchase
- 24x7 Customer Support
- Average 99.9% Success Rate
- More than 800,000 Satisfied Customers Worldwide
- Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications. You can view Vendor list of All Certification Exams offered:

<https://www.lead4pass.com/allproducts>

Need Help

Please provide as much detail as possible so we can best assist you.
To update a previously submitted ticket:



	<p>One Year Free Update</p> <p>Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>		<p>Money Back Guarantee</p> <p>To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>		<p>Security & Privacy</p> <p>We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
--	--	--	--	--	--

Any charges made through this site will appear as Global Simulators Limited.
All trademarks are the property of their respective owners.
Copyright © lead4pass, All Rights Reserved.