DVA-C01 Q&As

AWS Certified Developer - Associate (DVA-C01)

Pass Amazon DVA-C01 Exam with 100% Guarantee

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

https://www.pass2lead.com/aws-certified-developer-associate.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

A company needs to secure its existing website running behind an Elastic Load Balancer. The website\\'s Amazon EC2 instances are CPU-constrained. What should be done to secure the website while not increasing the CPU load on the EC2 web servers? (Select TWO.)

A. Configure an Elastic Load Balancer with SSL pass-through.

B. Configure SSL certificates on an Elastic Load Balancer.

C. Configure an Elastic Load Balancer with a Loadable Storage System.

D. Install SSL certificates on the EC2 instances.

E. Configure an Elastic Load Balancer with SSL termination.

Correct Answer: BD

QUESTION 2

A developer wants to modify the following AWS Cloud Formation template to embed another CloudFormation stack: Which syntax should the developer add to the blank line of the CloudFormation template to meet this requirement?

```
{
 "AWSTemplateFormatVersion" : "2010-09-09",
  "Resources" : {
    "cfStack" : {
          "Properties" : {
        "TemplateURL" : "https://s3.amazonaws.com/cloudformatio
        "Parameters" : {
          "InstanceType" : "t3.small"
        }
      }
    }
  }
}
```

A. "Mapping" : "AWS::CloudFormation::Stack",

B. "Type" : "AWS;:CloudFcrmation::NestedStack",

C. "Typequot; :";AWS::CloudFormation::Stac";,

D. "Mapping" : "AWS::CloudFormation::NestedStack",

Correct Answer: A

QUESTION 3

A data-processing application includes an AWS Lambda function that processes data in several steps. Recently, the function has been reaching the Lambda tii A developer wants to use AWS X-Ray to find out how long each step is taking so that the developer can determine which step is causing the timeout.

Which combination of actions should the developer take to accomplish this goal? (Select TWO.)

A. Modify the application to call the PutMetricData API operation after each processing step. Include the time taken in milliseconds.

B. Use the aws lambda update-function-configuration AWS CLI command to enable active tracing on the Lambda function.

C. Modify the application to record each processing step in an X-Ray subsegment by using the X-Ray software development kit (SDK).

D. Add the xray:PutTraceSegments permission and the xray:PutTelemetryRecords permission to the Lambda function\\'s execution role.

E. Modify the application to put each processing step in a separate Lambda layer. Include all the layers in the Lambda function.

Correct Answer: B

QUESTION 4

A nightly batch job loads 1 million new records into a DynamoDB table. The records are only needed for one hour, and the table needs to be empty by the next night\\'s batch job. Which is the MOST efficient and cost-effective method to provide an empty table?

A. Use DeleteItem using a ConditionExpression.

B. Use BatchWriteItem to empty all of the rows.

C. With a recursive function that scans and calls out DeleteItem.

D. Create and then delete the table after the task has completed.

Correct Answer: D

"Deleting an entire table is significantly more efficient than removing items one-by-one, which essentially doubles the write throughput as you do as many delete operations as put operations"

QUESTION 5

A company deploys a local news website on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). The company wants to implement authentication for viewers through well-known social identity providers (IdPs). The company\\'s development team must use AWS native services in the solution.

Which combination of actions will meet these requirements? (Choose two.)

A. Configure Amazon Cognito user pools with social IdPs

B. Configure OpenID Connect (OIDC)-compliant IdP endpoints

C. Create an AWS Lambda authorizer as a target in the ALB target group

D. Configure an ALB listener to add a rule for authentication

E. Use Lambda@Edge to authorize viewer requests in the ALB

Correct Answer: BD

Reference https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_create_oidc.html

QUESTION 6

An application is running on a cluster of Amazon EC2 instances. While trying to read objects stored within a single Amazon S3 bucket that are encrypted with server-side encryption with AWS KMS managed keys (SSE-KMS), the application receives the following error:

```
Service: AWSKMS; Status Code: 400; Error Code: ThrottlingException
```

Which combination of steps should be taken to prevent this failure? (Choose two.)

A. Contact AWS Support to request an AWS KMS rate limit increase.

B. Perform error retries with exponential backoff in the application code.

C. Contact AWS Support to request a S3 rate limit increase.

D. Import a customer master key (CMK) with a larger key size.

E. Use more than one customer master key (CMK) to encrypt S3 data

Correct Answer: AD

QUESTION 7

A developer is building a web application. The application will read temperature information from an Amazon DynamoDB table and will display the information for users. Each record is 5 KB to 7 KB in size. The application can request up to 80 items each second. The application must always return the most recent temperature values from the table.

How much read throughput is required to meet this load?

A. 40 read capacity units (RCUs)

B. 80 read capacity units (RCUs)

C. 160 read capacity units (RCUs)

D. 560 read capacity units (RCUs)

Correct Answer: A

QUESTION 8

A developer has three microservice projects that are separated into different folders under the same AWS CodeCommit repository. Each project has a separate AWS CodePipeline pipeline. The developer notices that when changes are

pushed to one microservice, all three pipelines begin to run.

The developer needs to ensure that only relevant pipelines run. The developer cannot make any changes to how the repository is organized.

Which solution will meet these requirements?

A. For each of the three microservice projects, create a separate CodeCommit repository.

B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function to evaluate changes to the repository and run the appropriate pipeline.

C. Create an Amazon API Gateway API that is backed by an AWS Lambda function to determine the appropriate pipeline to run. Add the API endpoint to a webhook in CodeCommit.

D. Migrate all three pipelines to a single pipeline. Add conditional stages to build a certain microservice project.

Correct Answer: C

Reference: https://aws.amazon.com/blogs/devops/adding-custom-logic-to-aws-codepipeline-with-aws-lambda-and-amazon-cloudwatch-events/

QUESTION 9

If an application is storing hourly log files from thousands of instances from a high traffic web site, which naming scheme would give optimal performance on S3?

A. Sequential

B. instanceID_log-HH-DD-MM-YYYY

C. instanceID_log-YYYY-MM-DD-HH

D. HH-DD-MM-YYYY-log_instanceID

E. YYYY-MM-DD-HH-log_instanceID

Correct Answer: B

Reference: https://acloud.guru/forums/aws-certified-developer-associate/discussion/-KU2dEtJb-LI5lSbH_S4/if- an-application-is-storing-hourly-log-files-from-thousands-of-instances-from-a

QUESTION 10

A developer has created a data collection application that uses Amazon API Gateway, AWS Lambda, and Amazon S3. The application\\'s users periodically upload data files and wait for the validation status to be reflected on a processing dashboard. The validation process is complex and time-consuming for large files.

Some users are uploading dozens of large files and have to wait and refresh the processing dashboard to see if the files have been validated. The developer must refactor the application to immediately update the validation result on the user\\'s dashboard without reloading the full dashboard.

What is the MOST operationally efficient solution that meets these requirements?

A. Integrate the client with an API Gateway WebSocket API. Save the user-uploaded files with the WebSocket connection ID. Push the validation status to the connection ID when the processing is complete to initiate an update of the user interface.

B. Launch an Amazon EC2 micro instance, and set up a WebSocket server. Send the user-uploaded file and user detail to the EC2 instance after the user uploads the file. Use the WebSocket server to send updates to the user interface when the uploaded file is processed.

C. Save the user\\'s email address along with the user-uploaded file. When the validation process is complete, send an email notification through Amazon Simple Notification Service (Amazon SNS) to the user who uploaded the file.

D. Save the user-uploaded file and user detail to Amazon DynamoDB. Use Amazon DynamoDB Streams with Amazon Simple Notification Service (Amazon SNS) push notifications to send updates to the browser to update the user interface.

Correct Answer: C

Reference: https://docs.aws.amazon.com/sns/latest/dg/sns-dg.pdf

QUESTION 11

A company is building an application to track athlete performance using an Amazon DynamoDB table. Each item in the table is identified by a partition key (user_id) and a sort key (sport_name). The table design is shown below:

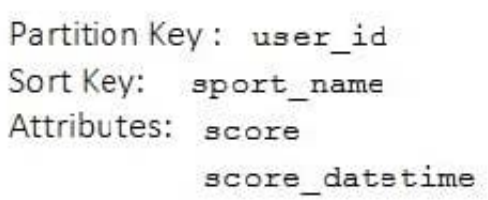

(Note: Not all table attributes are shown)

A Developer is asked to write a leaderboard application to display the top performers (user_id) based on the score for each sport_name.

What process will allow the Developer to extract results MOST efficiently from the DynamoDB table?

A. Use a DynamoDB query operation with the key attributes of user_id and sport_name and order the results based on the score attribute.

B. Create a global secondary index with a partition key of sport_name and a sort key of score, and get the results

C. Use a DynamoDB scan operation to retrieve scores and user_id based on sport_name, and order the results based on the score attribute.

D. Create a local secondary index with a primary key of sport_name and a sort key of score and get the results based on the score attribute.

Correct Answer: B

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html
https://docs.aws.amazon.com/zh_cn/amazondynamodb/latest/developerguide/GSI.html

QUESTION 12

A Developer is asked to implement a caching layer in front of Amazon RDS. Cached content is expensive to regenerate in case of service failure. Which implementation below would work while maintaining maximum uptime?

A. Implement Amazon ElastiCache Redis in Cluster Mode

B. Install Redis on an Amazon EC2 instance.

C. Implement Amazon ElastiCache Memcached.

D. Migrate the database to Amazon Redshift.

Correct Answer: A

https://docs.aws.amazon.com/AmazonElastiCache/latest/mem-ug/SelectEngine.html

QUESTION 13

A company maintains an application responsible for processing several thousand external callbacks each day. The company\\'s System administrators want to know how many callbacks are being received on a rolling basis, and they want this data available for 10 days. The company also wants the ability to issue automated alerts if the number of callbacks exceeds the defined thresholds.

What is the MOST cost-effective way to address the need to track and alert on these statistics?

A. Push callback data to an Amazon RDS database that can be queried to show historical data and to alert on exceeded thresholds.

B. Push callback data to AWS X-Ray and use AWS Lambda to query, display, and alert on exceeded thresholds.

C. Push callback data to Amazon Kinesis Data Streams and invoke an AWS Lambda function that stores data in Amazon DynamoDB and sends the required alerts.

D. Push callback data to Amazon CloudWatch as a custom metric and use the CloudWatch alerting mechanisms to alert System Administrators.

Correct Answer: D

QUESTION 14

Queries to an Amazon DynamoDB table are consuming a large amount of read capacity. The table has a significant number of large attributes. The application does not need all of the attribute data. How can DynamoDB costs be minimized while maximizing application performance?

A. Batch all the writes, and perform the write operations when no or few reads are being performed.

B. Create a global secondary index with a minimum set of projected attributes.

C. Implement exponential backoffs in the application.

D. Load balance the reads to the table using an Application Load Balancer.

Correct Answer: B

https://docs.aws.amazon.com/AWSEC2/latest/APIReference/query-api-troubleshooting.html

QUESTION 15

A company is adding items to an Amazon DynamoDB table from an AWS Lambda function that is written in Python. A developer needs to implement a solution that inserts records in the DynamoDB table and performs automatic retry when the insert fails.

Which solution meets these requirements with MINIMUM code changes?

A. Configure the Python code to run the AWS CLI through shell to call the PutItem operation

B. Call the PutItem operation from Python by using the DynamoDB HTTP API

C. Queue the items in AWS Glue, which will put them into the DynamoDB table

D. Use the AWS software development kit (SDK) for Python (boto3) to call the PutItem operation

Correct Answer: D

Reference: https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GettingStarted.Python.html