


DumpCollection

IT Exam Training online / Bootcamp



<http://www.dumpcollection.com>

PDF and Testing Engine, study and practice

Exam : **DVA-C01**

Title : AWS Certified Developer
Associate Exam

Vendor : Amazon

Version : DEMO

NO.1 A developer is building an application integrating an Amazon API Gateway with an AWS Lambda function. When calling the API, the developer receives the following error. Wed Nov 03 01:13:00 UTC 2017 : Method completed with status: 502 What should the developer do to resolve the error?

- A. Change the format of the payload sent to the API Gateway.
- B. Change the authorization header in the API call to access the Lambda function.
- C. Change the format of the Lambda function response to the API call.
- D. Change the HTTP endpoint of the API to an HTTPS endpoint.

Answer: C

NO.2 A company is building a serverless application that uses AWS Lambda
a. The application includes Lambda functions that are exposed by Amazon API Gateway The functions will use several large third-party libraries, and the build artifacts will exceed 50 MB in size. Which combination of steps should a developer take to prepare and perform the deployment? (Select TWO.)

- A. Issue the aws cloudformation package CLI command.
- B. Issue the aws lambda update-function-code CLI command with the -s3-bucket and -s3-key parameters.
- C. Issue the aws lambda update-function-code CLI command with a parameter that points to the source code in AWS CodeCommit.
- D. Issue the aws lambda update-function-code CLI command with the -zip-file fileb://my-function.zip parameter
- E. Upload the build artifact to Amazon S3.

Answer: B,E

NO.3 An AWS Lambda function accesses two Amazon DynamoDB tables. A developer wants to improve the performance of the Lambda function by identifying bottlenecks in the function. How can the developer inspect the timing of the DynamoDB API calls?

- A. Limit Lambda to no more than five concurrent invocations Monitor from the Lambda console
- B. Enable AWS X-Ray tracing for the function. View the traces from the X-Ray service.
- C. Add DynamoDB as an event source to the Lambda function. View the performance with Amazon CloudWatch metrics
- D. Place an Application Load Balancer (ALB) in front of the two DynamoDB tables. Inspect the ALB logs

Answer: B

NO.4 A developer wants to secure sensitive configuration data such as passwords, database strings, and application license codes. Access to this sensitive information must be tracked for future audit purposes.

- A. In an encrypted file on the source code bundle; grant the application access with Amazon IAM
- B. In the Amazon EC2 Systems Manager Parameter Store; grant the application access with IAM
- C. As an object in on Amazon S3 bucket, grant on Amazon EC2 instance access with on IAM rob.
- D. On an Amazon EBS encrypted volume attach the volume to an Amazon EC2 instance to access the data

Answer: B

NO.5 Which features can be used to restrict access to data in S3? Choose 2 answers

- A. Set an S3 Bucket policy.
- B. Use S3 Virtual Hosting
- C. Set an S3 ACL on the bucket or the object.
- D. Enable IAM Identity Federation.
- E. Create a CloudFront distribution for the bucket

Answer: A,C

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/secure-s3-resources/>

NO.6 A company uses Amazon DynamoDB for managing and tracking orders. The DynamoDB table is partitioned based on the order date. The company receives a huge increase in orders during a sales event, causing DynamoDB writes to throttle, and the consumed throughput is far below the provisioned throughput.

According to AWS best practices, how can this issue be resolved with MINIMAL costs?

- A. Create a new DynamoDB table for every order date.
- B. Increase the read and write capacity units of the DynamoDB table.
- C. Add a global secondary index to the DynamoDB table.
- D. Add a random number suffix to the partition key values.

Answer: D

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-partition-key-uniform-load.html>

NO.7 A company is using continuous integration and continuous delivery systems. A Developer now needs to automate a software package deployment to both Amazon EC2 instances and virtual servers running on-premises.

Which AWS service should be used to accomplish this?

- A. AWS CodePipeline
- B. AWS Elastic Beanstalk
- C. AWS CodeDeploy
- D. AWS CodeBuild

Answer: C

NO.8 A Developer must re-implement the business logic for an order fulfilment system. The business logic has to make requests to multiple vendors to decide where to purchase an item. The whole process can take up to a week to complete.

What is the MOST efficient and SIMPLEST way to implement a system that meets these requirements?

- A. Use AWS Lambda to asynchronously call a Lambda function for each vendor, and join the results.
- B. Use Amazon CloudWatch Events to orchestrate the Lambda functions.
- C. Use AWS Step Functions to execute parallel Lambda functions, and join the results.

D. Create an AWS SQS for each vendor, poll the queue from a worker instance, and joint the results.

Answer: C

Explanation:

<https://aws.amazon.com/step-functions/>

NO.9 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 scan operation to retrieve scores and user_id based on sport_name, and order the results based on the score attribute.

B. 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.

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

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

Answer: D

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html>

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

NO.10 A company maintains a REST service using Amazon API Gateway and the API Gateway native API key validation. The company recently launched a new registration page, which allows users to sign up for the service. The registration page creates a new API key using CreateApiKey and sends the new key to the user. When the user attempts to call the API using this key, the user receives a 403 Forbidden error. Existing users are unaffected and can still call the API.

What code updates will grant these new users access to the API?

A. The importApiKeys method must be called to import all newly created API keys into the current stage of the API.

B. The createUsagePlanKey method must be called to associate the newly created API key with the correct usage plan.

C. The createDeployment method must be called so the API can be redeployed to include the newly created API key.

D. The updateAuthorizer method must be called to update the API's authorizer to include the newly created API key.

Answer: B

Explanation:

<https://stackoverflow.com/questions/39061041/using-an-api-key-in-amazon-api-gateway>

NO.11 A company is launching a polling application. The application will store the results of each pool an Amazon DynamoDB table. Management wants to remove pool data after a few data and

store an archive of those records in Amazon S3.

Which approach would allow the application to archive each poll's data while keeping complexity to a MINIMUM?

- A. Enable cross-Region replication on the S3 bucket to achieve the poll data.
- B. Schedule an AWS Lambda function to periodically scan the DynamoDB table. Use the BatchWritten operation to delete the results of a scan Enable DynamoDB Stream on the table and store the records removed from the stream in Amazon S3.
- C. Enable DynamoDB Streams on the table. Configure the stream as trigger for AWS Lambda. Save records to Amazon S3 when records on the stream are modified.
- D. Enable Time to Live (TTL) on the DynamoDB table. Enable DynamoDB Streams on the table and store the records removed from the stream in Amazon S3.

Answer: C

NO.12 After installing the AWS CLI, a Developer tries to run the command `aws configure` but receives the following error:

Error: aws: command not found

What is the most likely cause of this error?

- A. Incorrect AWS credentials were provided.
- B. The aws executable is not in the PATH environment variable.
- C. Access to the aws executable has been denied to the installer.
- D. The aws script does not have an executable file mode.

Answer: B

Explanation:

<https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-troubleshooting.html>

NO.13 A developer is building an application using an Amazon API Gateway REST API backed by an AWS Lambda function that interacts with an Amazon DynamoDB table During testing, the developer observes high latency when making requests to the API How can the developer evaluate the end-to-end latency and identify performance bottlenecks?

- A. Enable and configure AWS X-Ray tracing on API Gateway and the Lambda function Use X-Ray to trace and analyze user requests
- B. Enable VPC Flow Logs to capture and analyze network traffic within the VPC
- C. Enable AWS CloudTrail logging and use the logs to map each latency and bottleneck
- D. Enable Amazon CloudWatch Logs for the Lambda function Enable execution logs for API Gateway to view and analyze user request logs.

Answer: A

NO.14 A Developer is building a mobile application and needs any update to user profile data to be pushed to all devices accessing the specific identity. The Developer does not want to manage a back end to maintain the user profile data.

What is the MOST efficient way for the Developer to achieve these requirements using Amazon Cognito?

- A. Use Cognito events.
- B. Use Cognito federated identities.

C. Use a Cognito user pool.

D. Use Cognito Sync.

Answer: D

Explanation:

Amazon Cognito Sync is an AWS service and client library that enables cross-device syncing of application-related user data. You can use it to synchronize user profile data across mobile devices and the web without requiring your own backend.

<https://docs.aws.amazon.com/cognito/latest/developerguide/cognito-sync.html>

NO.15 A Developer is creating an Auto Scaling group whose instances need to publish a custom metric to Amazon CloudWatch.

Which method would be the MOST secure way to authenticate a CloudWatch PUT request?

A. Modify the CloudWatch metric policies to allow the PutMetricData permission to instances from the Auto Scaling group.

B. Create an IAM user with PutMetricData permission and put the user credentials in a private repository; have applications pull the credentials as needed.

C. Create an IAM user with PutMetricData permission, and modify the Auto Scaling launch configuration to inject the user credentials into the instance user data.

D. Create an IAM role with PutMetricData permission and modify the Auto Scaling launching configuration to launch instances using that role.

Answer: D

NO.16 A developer is leveraging a Border Gateway Protocol (BGP)-based AWS VPN connection to connect from on-premises to Amazon EC2 instances in the developer's account. The developer is able to access an EC2 instance in subnet A, but is unable to access an EC2 instance in subnet B in the same VPC. Which logs can the developer use to verify whether the traffic is reaching subnet B?

A. AWS CloudTrail logs

B. VPC Flow Logs

C. BGP logs

D. VPN logs

Answer: B

NO.17 A Developer is storing sensitive documents in Amazon S3 that will require encryption at rest. The encryption keys must be rotated annually, at least.

What is the easiest way to achieve this?

A. Import a custom key into AWS KMS with annual rotation enabled

B. Export a key from AWS KMS to encrypt the data

C. Use AWS KMS with automatic key rotation

D. Encrypt the data before sending it to Amazon S3

Answer: C

Explanation:

<https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html>

<https://docs.aws.amazon.com/kms/latest/developerguide/custom-key-store-overview.html> You can use the same techniques to view and manage the CMKs in your custom key store that you use for

CMKs in the AWS KMS key store. You can control access with IAM and key policies, create tags and aliases, enable and disable the CMKs, and schedule key deletion. You can use the CMKs for cryptographic operations and use them with AWS services that integrate with AWS KMS. However, you cannot enable automatic key rotation and you cannot import key material into a CMK in a custom key store.

Q: Can I rotate my keys? Yes. You can choose to have AWS KMS automatically rotate CMKs every year, provided that those keys were generated within AWS KMS HSMs. Automatic key rotation is not supported for imported keys, asymmetric keys, or keys generated in an AWS CloudHSM cluster using the AWS KMS custom key store feature. If you choose to import keys to AWS KMS or asymmetric keys or use a custom key store, you can manually rotate them by creating a new CMK and mapping an existing key alias from the old CMK to the new CMK. <https://aws.amazon.com/kms/faqs/>

NO.18 A developer wants to migrate a Windows .NET application that is running on IIS with a Microsoft SQL Server database to AWS. The developer does not want to think about provisioning and managing the infrastructure.

What should the developer do to migrate the application with the LEAST amount of effort?

- A.** Launch Amazon EC2 instances for Windows Server. Back up and restore the database to Amazon RDS. Deploy the web application to the new EC2 instances
- B.** Migrate the database to Amazon DynamoDB Use Amazon API Gateway and AWS Lambda to create a web application interface that is hosted in an Amazon S3 bucket.
- C.** Back up and restore the database to Amazon RDS. Use the .NET Migration Assistant for AWS Elastic Beanstalk to migrate the web application to a preconfigured solution stack that Elastic Beanstalk provides.
- D.** Containerize the application on premises. Push the image to Amazon Elastic Container Registry (Amazon ECR). Create an AWS CloudFormation template to

Answer: C

deploy the application

NO.19 A developer is building an AWS Lambda function that will dynamically generate and send a weekly newsletter to 100,000 users. This newsletter contains both static text and images. The developer needs a fast and highly scalable place to store the images that will be hyperlinked in the newsletter. Where should the developer store these images?

- A.** Use an Amazon S3 backed Amazon CloudFront distribution with a high Time-to-Live (TTL) to maximize caching
- B.** Use an Amazon S3 bucket and S3 Transfer Acceleration to speed up the image download
- C.** Use an Amazon DynamoDB table with DynamoDB Streams and read capacity auto scaling enabled
- D.** Use an Amazon Aurora database with a public DNS endpoint and auto scaling enabled

Answer: A

NO.20 A developer is writing an application that will run on Amazon EC2 instances in an Auto Scaling group. The developer wants to externalize the session state to support the application. Which AWS services or resources can the developer use to meet these requirements? (Select TWO.)

- A.** Amazon Simple Queue Service (Amazon SQS)
- B.** Application Load Balancer
- C.** Amazon DynamoDB

D. Amazon ElastiCache

E. Amazon Cognito

Answer: C,D

NO.21 A large company has its application components distributed across.. company needs to collect and visualize trace data across these accounts.

What should be used to meet these requirements?

A. Amazon Elasticsearch Service

B. Amazon VPC flow logs

C. Amazon CloudWatch

D. AWS X-Ray

Answer: D

NO.22 A company is creating a continuous integration and continuous delivery (CI/CD) process by using AWS CodePipeline for its application on AWS The CI/CD process will pull code from an AWS CodeCommit repository, create the application infrastructure by using AWS Cloud Formation, deploy the frontend code to an Amazon S3 bucket that is configured for static website hosting, and deploy the application backend on an Amazon Elastic Container Service (Amazon ECS) cluster.

A developer needs to create a new CodePipeline stage that creates the application infrastructure.

Which solution will meet these requirements with the LEAST operational overhead?

A. Create a new action with AWS CodeBuild as the action provider Configure the buildspec to make an API call by using the AWS CLI to create the CloudFormation stack

B. Create a new action with AWS Lambda as the action provider Create a Lambda function that makes an AWS SDK API call to create the CloudFormation stack

C. Create a new action with CloudFormation as the action provider Set the action mode to CREATE UPDATE Target the CloudFormation stack to be launched.

D. Create a new action with Jenkins as the action provider. Create and configure a Jenkins job to make an API call by using the AWS CLI to create the CloudFormation stack.

Answer: A

NO.23 An existing serverless application processes uploaded image files. The process currently uses a single Lambda function that takes an image file, performs the processing, and stores the file in Amazon S3. Users of the application now require thumbnail generation of the images. Users want to avoid any impact to the time it takes to perform the image uploads.

How can thumbnail generation be added to the application, meeting user requirements while minimizing changes to existing code?

A. Create a second Lambda function that handles thumbnail generation and storage. Change the existing Lambda function to invoke it asynchronously.

B. Create an S3 event notification to an SQS Queue. Create a scheduled Lambda function that processes the queue, and generates and stores thumbnails.

C. Create an S3 event notification with a Lambda function destination. Create a new Lambda function to generate and store thumbnails.

D. Change the existing Lambda function handling the uploads to create thumbnails at the time of upload. Have the function store both the image and thumbnail in Amazon S3.

Answer: C

Explanation:

<https://docs.aws.amazon.com/lambda/latest/dg/with-s3-example.html>

NO.24 A Developer has an application that can upload tens of thousands of objects per second to Amazon S3 in parallel within a single AWS account. As part of new requirements, data stored in S3 must use server side encryption with AWS KMS (SSE-KMS). After creating this change, performance of the application is slower.

Which of the following is MOST likely the cause of the application latency?

- A.** The client encryption of the objects is using a poor algorithm.
- B.** KMS requires that an alias be used to create an independent display name that can be mapped to a CMK.
- C.** Amazon S3 throttles the rate at which uploaded objects can be encrypted using Customer Master Keys.
- D.** The AWS KMS API calls limit is less than needed to achieve the desired performance.

Answer: D

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2018/08/aws-key-management-service-increases-api-requests-per-second-limits/> KMS API access limit is 10k/sec in us-east and some others and 5.5k/sec for the rest of the regions. Client can request this limit to be changed.

NO.25 A company is developing an application that will run on several Amazon EC2 instances in an Auto Scaling group and can access a database running on Amazon EC2. The application needs to store secrets required to connect to the database. The application must allow for periodic secret rotation, and there should be no changes to the application when a secret changes.

What is the SAFEST way to meet these requirements?

- A.** Associate an IAM role to the EC2 instance where the application is running with permission to access the database.
- B.** Use AWS Systems Manager Parameter Store with the SecureString data type to store secrets.
- C.** Hard code the database secrets in the application code itself.
- D.** Configure the application to store secrets in Amazon S3 object metadata.

Answer: B

NO.26 A developer deployed an application to an Amazon EC2 instance. The application needs to know the public IPv4 address of the instance. How can the application find this information?

- A.** Query the instance metadata from <http://169.254.169.254/latest/meta-data/>.
- B.** Query the Amazon Machine Image (AMI) information from <http://169.254.169.254/latest/meta-data/ami/>.
- C.** Check the hosts file of the operating system.
- D.** Query the instance user data from <http://169.254.169.254/latest/user-data/>.

Answer: A

NO.27 A developer must build a mobile application that allows users to read and write data from an Amazon DynamoDB table to store user state for each unique user. The solution needs to limit data

access to allow users access only to their own data Which solution below is the most secure?

- A. Modify the DynamoDB table to allow public read and writes, then add client-side filtering
- B. Create a web portal for users to create an account on AWS Directory Service
- C. Use Amazon Cognito identity pools to assign unique identifiers and provide user access
- D. Embed AWS access credentials into the application and create DynamoDB queries that limit user access.

Answer: A

NO.28 A developer is building an application that runs behind an application Load Balancer (ALB). The application is configured as the origin for an Amazon CloudFront distribution. Users will log in to the application using their social media accounts.

How can the developer authenticate and authorize users?

- A. Authorize the users by calling the Amazon Cognito API in the AWS Lambda authorizer on the ALB
- B. Configure Cloudfront to use Amazon Cognito as one of the authentication providers
- C. Configure the ALB to use Amazon Cognito as one of the authentication providers
- D. Validate the user by inspecting the tokens using AWS Lambda authorizers on the ALB

Answer: B

NO.29 A company is building a compute-intensive application that will run on a fleet of Amazon EC2 instances. The application uses attached Amazon EBS disks for storing data. The application will process sensitive information and all the data must be encrypted.

What should a developer do to ensure the data is encrypted on disk without impacting performance?

- A. Add logic to write all data to an encrypted Amazon S3 bucket
- B. Create a new Amazon Machine Image (AMI) with an encrypted root volume and store the data to ephemeral disks.
- C. Configure the Amazon EC2 instance fleet to use encrypted EBS volumes for storing data
- D. Add a custom encryption algorithm to the application that will encrypt and decrypt all data

Answer: C

NO.30 A developer is working on a serverless application. The application uses Amazon API Gateway, AWS Lambda functions that are written in Python, and Amazon DynamoDB.

Which combination of steps should the developer take so that the Lambda functions can be debugged in the event of application failures? (Select TWO)

- A. Ensure that the execution role for the Lambda function has access to write to Amazon CloudWatch Logs.
- B. Configure an AWS CloudTrail trail to deliver log files to an Amazon S3 bucket
- C. Enable an AWS CloudTrail trail for the Lambda function
- D. Ensure that the Lambda functions write log messages to stdout and stderr
- E. Use the Amazon CloudWatch metric for Lambda errors to create a CloudWatch alarm.

Answer: A,E

NO.31 A company has a multi-tiered web application on AWS. During a recent spike in traffic, one of the primary relational databases on Amazon RDS could not serve all the traffic. Some read queries for repeatedly accessed items failed, so users received error messages.

What can be done to minimize the impact on database read queries MOST efficiently during future traffic spikes?

- A. Use Amazon S3 to cache database query results.
- B. Use local storage and memory on Amazon EC2 instances to cache data.
- C. Use Amazon RDS as a custom origin for Amazon CloudFront.
- D. Use Amazon ElastiCache in front of the primary database to cache data.

Answer: D

NO.32 A developer is planning to use an Amazon API Gateway and AWS Lambda to provide a REST API. The developer will have three distinct environments to manage development, test, and production. How should the application be deployed while minimizing the number of resources to manage?

- A. Create one API Gateway with multiple stages with one Lambda function with multiple aliases.
- B. Assign a Region for each environment and deploy API Gateway and Lambda to each Region.
- C. Create one API Gateway and one Lambda function, and use a REST parameter to identify the environment.
- D. Create a separate API Gateway and separate Lambda function for each environment in the same Region.

Answer: A

NO.33 An application writes items to an Amazon DynamoDB table. As the application scales to thousands of instances, calls to the DynamoDB API generate occasional ThrottlingException errors. The application is coded in a language incompatible with the AWS SDK. How should the error be handled?

- A. Use Amazon SQS as an API message bus.
- B. Add exponential backoff to the application logic.
- C. Send the items to DynamoDB through Amazon Kinesis Data Firehose.
- D. Pass API calls through Amazon API Gateway.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/throttled-ddb/>

SDKs automatically add exponential backoff. If not using the AWS SDKs, add your own backoff logic to the application code.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Programming.Errors.html#Programming.Errors.Handling>

NO.34 An application runs on multiple EC2 instances behind an ELB.

Where is the session data best written so that it can be served reliably across multiple requests?

- A. Write data to Amazon ElasticCache.
- B. Write data to the root filesystem.
- C. Write data to Amazon Elastic Block Store.
- D. Write data to Amazon EC2 instance Block Store.

Answer: A

NO.35 Which statements about DynamoDB are true? Choose 2 answers

- A.** DynamoDB uses optimistic concurrency control
- B.** DynamoDB uses a pessimistic locking model
- C.** DynamoDB uses conditional writes for consistency
- D.** DynamoDB restricts item access during writes
- E.** DynamoDB restricts item access during reads

Answer: A,C