



- Expert Verified, Online, **Free**.

A company is implementing an application on Amazon EC2 instances. The application needs to process incoming transactions. When the application detects a transaction that is not valid, the application must send a chat message to the company's support team. To send the message, the application needs to retrieve the access token to authenticate by using the chat API.


A developer needs to implement a solution to store the access token. The access token must be encrypted at rest and in transit. The access token must also be accessible from other AWS accounts.

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

- A. Use an AWS Systems Manager Parameter Store SecureString parameter that uses an AWS Key Management Service (AWS KMS) AWS managed key to store the access token. Add a resource-based policy to the parameter to allow access from other accounts. Update the IAM role of the EC2 instances with permissions to access Parameter Store. Retrieve the token from Parameter Store with the decrypt flag enabled. Use the decrypted access token to send the message to the chat.
- B. Encrypt the access token by using an AWS Key Management Service (AWS KMS) customer managed key. Store the access token in an Amazon DynamoDB table. Update the IAM role of the EC2 instances with permissions to access DynamoDB and AWS KMS. Retrieve the token from DynamoDB. Decrypt the token by using AWS KMS on the EC2 instances. Use the decrypted access token to send the message to the chat.
- C. Use AWS Secrets Manager with an AWS Key Management Service (AWS KMS) customer managed key to store the access token. Add a resource-based policy to the secret to allow access from other accounts. Update the IAM role of the EC2 instances with permissions to access Secrets Manager. Retrieve the token from Secrets Manager. Use the decrypted access token to send the message to the chat.
- D. Encrypt the access token by using an AWS Key Management Service (AWS KMS) AWS managed key. Store the access token in an Amazon S3 bucket. Add a bucket policy to the S3 bucket to allow access from other accounts. Update the IAM role of the EC2 instances with permissions to access Amazon S3 and AWS KMS. Retrieve the token from the S3 bucket. Decrypt the token by using AWS KMS on the EC2 instances. Use the decrypted access token to send the message to the chat.

Correct Answer: C

Community vote distribution



Untamables Highly Voted 1 year, 10 months ago

Selected Answer: C

The correct answer is C.

<https://aws.amazon.com/premiumsupport/knowledge-center/secrets-manager-share-between-accounts/>

https://docs.aws.amazon.com/secretsmanager/latest/userguide/auth-and-access_examples_cross.html

Option A is wrong. It seems to be a good solution. However, AWS managed keys cannot be used for cross account accessing.

upvoted 30 times

CyberBaby803 1 year, 10 months ago

Based on this AWS managed keys can be used for cross account accessing. <https://docs.aws.amazon.com/kms/latest/developerguide/key-policy-modifying-external-accounts.html>

upvoted 2 times

AgboolaKun 1 year, 7 months ago

I am not sure if the documentation you provided specifically say that AWS managed keys can be used for cross account accessing.

However, @Untamables explanation is on point. Please see this Stack Overflow thread -

<https://stackoverflow.com/questions/63420732/sharing-an-aws-managed-kms-key-with-another-account>

upvoted 1 times

jipark 1 year, 6 months ago

cross account, rotate is key for 'Security Manager'

upvoted 2 times

sumanshu Most Recent 1 month, 1 week ago

Selected Answer: C

For cross-account access, the AWS managed key (Option A) will be difficult to manage because it doesn't allow you to directly manage cross-account access. Therefore, Option C (AWS Secrets Manager with a customer-managed KMS key) is the recommended solution for cross-account

access and security.

upvoted 1 times

🗨️ **sumanshu** 1 month, 1 week ago

B) Eliminated - Creating and managing the DynamoDB table create overhead

upvoted 1 times

🗨️ **sumanshu** 1 month, 1 week ago

D) Eliminated - Complexity and overhead for Managing an S3 bucket with access policies for cross-account access

upvoted 1 times

🗨️ **trieudo** 1 month, 3 weeks ago

Selected Answer: C

keywords: LEAST management overhead

==> Discard B, D: you must do many steps to config with Storage, DB with KMS, IAM Role

==> Discard A: Pretty correct, but in use, you may write some script. It can work but requires additional configuration and doesn't offer some of the benefits tailored for secrets management like automatic rotation.

D: The solution with AWS Secrets Manager (option C) provides the least management overhead because:

Secrets Manager is specifically designed for storing and managing sensitive information like access tokens.

It natively integrates with AWS KMS for encryption and decryption.

It simplifies access control and auditing.

By adding a resource-based policy, cross-account access is easily managed without the need for additional configurations like DynamoDB tables or S3 bucket policies.

upvoted 1 times

🗨️ **trieudo** 1 month, 3 weeks ago

D: ==> C: ... sr for inconsistent

upvoted 1 times

🗨️ **Tee400** 4 months, 1 week ago

Selected Answer: A

Use an AWS Systems Manager Parameter Store SecureString parameter that uses an AWS Key Management Service (AWS KMS) AWS managed key to store the access token. This option allows you to securely store the access token in the Parameter Store, which automatically encrypts the data at rest and in transit. By adding a resource-based policy, you can also grant access to the access token from other AWS accounts. The IAM role of the EC2 instances can be updated to allow permissions to access the Parameter Store, and the access token can be retrieved with the decrypt flag enabled for use in sending the chat message. This option requires minimal setup and management compared to the other choices.

upvoted 1 times

🗨️ **cgpt** 4 months, 1 week ago

Selected Answer: A

By default, AWS Systems Manager Parameter Store does not natively support cross-account access for SecureString parameters. However, you can configure cross-account access to SecureString parameters by sharing the KMS key with the target AWS accounts. To do this, you need to create a resource-based KMS key policy that allows access to the key by the external AWS account(s). After configuring the KMS key policy to allow the necessary cross-account access, you can grant IAM roles in the target accounts permission to access the SecureString parameters that are encrypted using that KMS key.

upvoted 2 times

🗨️ **tomchandler077** 4 months, 1 week ago

AWS Secrets Manager (Option C) is designed for exactly this kind of use case, providing built-in functionality for secure storage and retrieval of secrets with minimal management overhead, especially for managing access tokens and cross-account access.

Amazon S3 with KMS (Option D), while familiar and powerful, requires more manual work to set up and manage the security aspects, which can lead to increased overhead in comparison to Secrets Manager.

Given that the goal is to have the least management overhead, Option C is the best fit because it is purpose-built for managing secrets and automates much of the complexity involved in secure storage and retrieval.

upvoted 1 times

🗨️ **Anandesh** 6 months, 2 weeks ago

Selected Answer: C

<https://docs.aws.amazon.com/secretsmanager/latest/userguide/data-protection.html>

https://docs.aws.amazon.com/secretsmanager/latest/userguide/auth-and-access_resource-policies.html

<https://docs.aws.amazon.com/secretsmanager/latest/userguide/security-encryption.html>

upvoted 2 times

🗨️ **nkroker** 7 months ago

C is the correct answer as the Secrets Manager supports resource-based policies, allowing you to grant access to other AWS accounts easily.

upvoted 1 times

🗨️ **NagaoShingo** 8 months ago

Selected Answer: C

C is correct answer.

upvoted 1 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗨️ **shabeebocoder** 10 months ago

Selected Answer: C

This is the correct answer for lease overhead to manage the secret key

upvoted 1 times

🗨️ **ibratoev** 10 months, 1 week ago

Selected Answer: C

It is C

upvoted 1 times

🗨️ **certplan** 10 months, 3 weeks ago

- Option A involves using AWS Systems Manager Parameter Store, which can work but requires additional configuration and doesn't offer some of the benefits tailored for secrets management like automatic rotation.

- Option B involves storing the access token in DynamoDB, which is not specifically designed for secrets management, and managing encryption and decryption manually using AWS KMS.

- Option D involves using S3, which again is not designed for secrets management, and adds complexity in managing access policies and permissions. Additionally, accessing the token would involve reading from S3, decrypting it, and then using it, which is less straightforward compared to using a service like Secrets Manager.

upvoted 2 times

🗨️ **SD_CS** 11 months, 3 weeks ago

Selected Answer: A

I think this would be A as this is cheaper than C. Any reason why A can not be the answer?

upvoted 1 times

🗨️ **TheFivePips** 11 months ago

From what I can find, You can apply resource-based policies at the Parameter Store level to control access to the entire Parameter Store service. However, you cannot apply resource-based policies directly to individual parameters within the Parameter Store.

That is seemingly the only reason I would choose C over A.

But were also not looking for whats cheapest, were looking for whats easiest to manage

upvoted 2 times

🗨️ **tsdsmth** 1 year ago

The answer would be C if an AWS-managed key was used, as Secrets Manager and KMS are good for situations like this. However, the use of a customer-managed key increases management overhead. So the best answer is D, not C.

upvoted 2 times

🗨️ **gilleep_17** 1 year ago

You cannot use a resource-based policy with a parameter in the Parameter Store. The stephen answer Option C is correct Practise paper3

upvoted 1 times

🗨️ **gilleep_17** 1 year ago

customer managed key , its an extra work. So I am confused with option A and C

upvoted 1 times

A company is running Amazon EC2 instances in multiple AWS accounts. A developer needs to implement an application that collects all the lifecycle events of the EC2 instances. The application needs to store the lifecycle events in a single Amazon Simple Queue Service (Amazon SQS) queue in the company's main AWS account for further processing.

Which solution will meet these requirements?

- A. Configure Amazon EC2 to deliver the EC2 instance lifecycle events from all accounts to the Amazon EventBridge event bus of the main account. Add an EventBridge rule to the event bus of the main account that matches all EC2 instance lifecycle events. Add the SQS queue as a target of the rule.
- B. Use the resource policies of the SQS queue in the main account to give each account permissions to write to that SQS queue. Add to the Amazon EventBridge event bus of each account an EventBridge rule that matches all EC2 instance lifecycle events. Add the SQS queue in the main account as a target of the rule.
- C. Write an AWS Lambda function that scans through all EC2 instances in the company accounts to detect EC2 instance lifecycle changes. Configure the Lambda function to write a notification message to the SQS queue in the main account if the function detects an EC2 instance lifecycle change. Add an Amazon EventBridge scheduled rule that invokes the Lambda function every minute.
- D. Configure the permissions on the main account event bus to receive events from all accounts. Create an Amazon EventBridge rule in each account to send all the EC2 instance lifecycle events to the main account event bus. Add an EventBridge rule to the main account event bus that matches all EC2 instance lifecycle events. Set the SQS queue as a target for the rule.

Correct Answer: D

Community vote distribution

D (81%)

Other

🗳️ 👤 **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: D

The correct answer is D.

Amazon EC2 instances can send the state-change notification events to Amazon EventBridge.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring-instance-state-changes.html>

Amazon EventBridge can send and receive events between event buses in AWS accounts.

<https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-cross-account.html>

upvoted 19 times

🗳️ 👤 **jipark** 1 year, 6 months ago

thanks a lot

upvoted 1 times

🗳️ 👤 **geekdamsel** Highly Voted 1 year, 8 months ago

This question came in exam. Correct answer is D.

upvoted 10 times

🗳️ 👤 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: D

https://docs.aws.amazon.com/us_en/eventbridge/latest/userguide/eb-cross-account.html#:~:text=You%20can%20configure%20EventBridge%20to%20send%20and%20receive,events%20from%20the%20event%20bus%20in%20your%

upvoted 1 times

🗳️ 👤 **sumanshu** 1 month, 1 week ago

A) Eliminated - Amazon EC2 itself does not send lifecycle events directly to EventBridge in another account. The EC2 lifecycle events would be published to EventBridge within the account where the EC2 instance resides.

The key here is using EventBridge cross-account event bus access. Amazon EventBridge allows for event buses to receive events from other AWS accounts if the appropriate permissions are set.

upvoted 1 times

🗳️ 👤 **sumanshu** 1 month, 1 week ago

B) Will work as Option B is implying a direct route from each AWS account's EventBridge to the main account's SQS queue, bypassing the main account's EventBridge event bus - Eliminated - as No centralized events, which is in Option D

upvoted 1 times

🗨️ **sumanshu** 1 month, 1 week ago

C) Eliminated - more maintenance overhead of LAMBDA

upvoted 1 times

🗨️ **trieudo** 1 month, 3 weeks ago

Selected Answer: D

==> Discard C: lambda scans ==> it will be delay by scan all data

==> Discard A: 'matches all EC2 instance' ==> hard to maintain, when updating many times can occurs

==> Discard B: it works, but it maybe have some security problem when pushing raw data (not clean) into SQS. It also doesn't take advantage of

D: By configuring the main account's event bus to accept events from other accounts and adding rules in those accounts to forward lifecycle events, this solution achieves secure and efficient centralization. EventBridge then routes the events to an SQS queue in the main account for further processing.

upvoted 2 times

🗨️ **Anandesh** 6 months, 2 weeks ago

Selected Answer: D

<https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-cross-account.html>

upvoted 1 times

🗨️ **NagaoShingo** 8 months ago

Selected Answer: D

D is correct answer.

upvoted 1 times

🗨️ **65703c1** 8 months, 2 weeks ago

D is the correct answer.

upvoted 1 times

🗨️ **xdkonorek2** 1 year, 1 month ago

Selected Answer: D

Tried to implement both B and D

It's tricky, because B could be possible but you can't select cross-account SQS as target to the rule, option D is 100% correct

upvoted 1 times

🗨️ **dongocanh272** 1 year, 2 months ago

Selected Answer: D

My answer is D

upvoted 2 times

🗨️ **Digo30sp** 1 year, 3 months ago

Selected Answer: D

Answer C is correct

upvoted 1 times

🗨️ **TeeTheMan** 1 year, 6 months ago

Selected Answer: B

Seems to me the correct answer is B. The current most voted answer is B, but can someone explain why it's better than B? I think B is better because it has fewer steps. The events go straight from each account into the queue. Unlike in D which has the intermediate step of the event bus of the main account. Also, why would you want to pollute the event bus of the main account with events from other accounts when it isn't necessary?

upvoted 5 times

🗨️ **KillThemWithKindness** 1 year, 6 months ago

B

Answer A is incorrect because Amazon EventBridge events can't be sent directly from one account's event bus to another.

Answer C is incorrect because it's unnecessary and inefficient to use Lambda to periodically scan all EC2 instances for lifecycle changes. Amazon EventBridge can capture these events automatically as they occur.

Answer D is incorrect because it is not possible to configure the main account event bus to receive events from all accounts directly, and Amazon

EventBridge events can't be sent directly from one account's event bus to another. The EventBridge rules need to be set up in the accounts where the events are generated.

upvoted 2 times

🗨️ 👤 **KillThemWithKindness** 1 year, 6 months ago

Sorry Im wrong, AWS allow to send and receive Amazon EventBridge events between AWS accounts.

<https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-cross-account.html>

Both B and D works, but D is more centralized

upvoted 4 times

🗨️ 👤 **ezredame** 1 year, 8 months ago

Selected Answer: D

The correct answer is D.

<https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-cross-account.html>

upvoted 2 times

🗨️ 👤 **Bibay** 1 year, 9 months ago

Selected Answer: A

Option D is not the best solution because it involves configuring the permissions on the main account's EventBridge event bus to receive events from all accounts, which can lead to potential security risks. Allowing other AWS accounts to send events to the main account's EventBridge event bus can potentially open up a security vulnerability, as it increases the attack surface area for the main account.

On the other hand, option A is the best solution because it involves using Amazon EventBridge, which is a serverless event bus that can be used to route events between AWS services or AWS accounts. By configuring Amazon EC2 to deliver the EC2 instance lifecycle events from all accounts to the Amazon EventBridge event bus of the main account, and adding the SQS queue as a target of the rule, the application can collect all the lifecycle events of the EC2 instances in a single queue in the main account without compromising the security posture of the AWS environment.

upvoted 1 times

🗨️ 👤 **ihebchorfi** 1 year, 9 months ago

Selected Answer: B

B solution meets all da requirements. By using resource policies, you can grant permissions for other accounts to write to the SQS queue in the main account.

Then, you create EventBridge rules in each account dat match EC2 lifecycle events and use da main account's SQS queue as a target for these rules. It's da best choice for dis scenario.

upvoted 1 times

🗨️ 👤 **MrTee** 1 year, 9 months ago

Selected Answer: D

This solution allows the collection of all the lifecycle events of the EC2 instances from multiple AWS accounts and stores them in a single Amazon SQS queue in the company's main AWS account for further processing

upvoted 2 times

🗨️ 👤 **shahs10** 1 year, 10 months ago

For Option C using lambda does not seem to be a good solution as we would have to trigger lambda on some schedule and it will has less granularity in time.

For D. Why would we be matching EC2 instance lifecycle events in Main account event bus and not in each account event bus and reducing overhead for main account

upvoted 1 times

An application is using Amazon Cognito user pools and identity pools for secure access. A developer wants to integrate the user-specific file upload and download features in the application with Amazon S3. The developer must ensure that the files are saved and retrieved in a secure manner and that users can access only their own files. The file sizes range from 3 KB to 300 MB.

Which option will meet these requirements with the HIGHEST level of security?

- A. Use S3 Event Notifications to validate the file upload and download requests and update the user interface (UI).
- B. Save the details of the uploaded files in a separate Amazon DynamoDB table. Filter the list of files in the user interface (UI) by comparing the current user ID with the user ID associated with the file in the table.
- C. Use Amazon API Gateway and an AWS Lambda function to upload and download files. Validate each request in the Lambda function before performing the requested operation.
- D. Use an IAM policy within the Amazon Cognito identity prefix to restrict users to use their own folders in Amazon S3.

Correct Answer: D

Community vote distribution

D (90%) 5%

Untamables Highly Voted 1 year, 10 months ago

Selected Answer: D

D

I actually apply this solution the production applications.

Examples

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_s3_cognito-bucket.html

<https://docs.amplify.aws/lib/storage/getting-started/q/platform/js/>

upvoted 11 times

sumanshu Most Recent 1 month, 1 week ago

Selected Answer: D

This solution provides the highest level of security by ensuring that each user can only access their own files in Amazon S3 based on their identity.

upvoted 1 times

sumanshu 1 month, 1 week ago

A) Eliminated - S3 Event Notifications are typically used for triggering processes when an object is uploaded or modified in S3, but they don't directly validate or enforce security controls on user-specific access

upvoted 1 times

sumanshu 1 month, 1 week ago

B) Eliminated - Storing file details in DynamoDB and using it to filter access in the UI could be helpful for organizing and tracking files.

However, this does not address the core security concern of restricting direct access to S3 objects themselves. Even though DynamoDB can help track files, it still doesn't enforce access control on the S3 objects themselves

upvoted 1 times

sumanshu 1 month, 1 week ago

C) Eliminated - Lambda would need to handle the file transfer and validation, which could introduce scaling issues for large file sizes (300 MB).

upvoted 1 times

trieudo 1 month, 3 weeks ago

Selected Answer: D

Highest secure



=> discard A: S3 Event Notifications, it act on notify with event, not blocking risk

=> discard B: You can't handle by UI or BE code by User ID, but hacker can try to access directly into storage s3

=> discard C: I would have said C but that is kind of a custom solution that is both more overhead and more prone to error

D ==> The identity prefix in Amazon Cognito allows you to create a unique identity for each user. This prefix can be used as part of the IAM policy to control access at a more granular level.

upvoted 2 times

  **Bibay** 4 months, 1 week ago

Selected Answer: C



D is not the best option as IAM policies only apply to actions taken through AWS Management Console, SDKs, and CLI. It does not apply to direct access to S3 from the application.

Option B is a good approach, but it requires additional overhead to manage the DynamoDB table.

Option A is also a possible solution but only provides limited security as it only validates the upload and download requests, and it does not provide user-level authorization.

Option C is the best choice as it allows the developer to implement a custom authentication mechanism in the Lambda function, providing the highest level of security. The authentication mechanism can be integrated with Amazon Cognito user pools and identity pools to authenticate users and ensure that only the owner of the file can upload and download it.

upvoted 1 times

  **grzess** 1 year, 8 months ago

Implementing custom authentication / authorization solution is extremely bad practice. Any developers is prone to mistakes. It's always better to trust the dedicated solution. Thus option C is definitely not the correct one.

upvoted 3 times

  **ACurryDeveloper** 6 months ago

D you benchods. It says cognito in the question motherchods



upvoted 1 times

  **NagaoShingo** 8 months ago

Selected Answer: D

D is correct answer.

upvoted 1 times

  **65703c1** 8 months, 2 weeks ago

D is the correct answer.

upvoted 1 times

  **TheFivePips** 11 months ago

Selected Answer: D

The identity prefix in Amazon Cognito allows you to create a unique identity for each user. This prefix can be used as part of the IAM policy to control access at a more granular level.

I would have said C but that is kind of a custom solution that is both more overhead and more prone to error

upvoted 2 times

  **tfmzworld** 1 year ago

Selected Answer: D

D is the answer

upvoted 1 times

  **Chimzi** 1 year ago

Selected Answer: D

B can work but does not provide the same level of security as D

upvoted 1 times

  **dongocanh272** 1 year, 2 months ago

Selected Answer: B

I consider between B & D

upvoted 1 times

  **Digo30sp** 1 year, 3 months ago

Selected Answer: D

Answer D is correct

upvoted 1 times

🗨️ 👤 **MrTee** 1 year, 9 months ago

Selected Answer: D

This solution ensures that users can access only their own files in a secure manner.

upvoted 3 times

🗨️ 👤 **haaris786** 1 year, 10 months ago

Answer D:

<https://docs.aws.amazon.com/cognito/latest/developerguide/amazon-cognito-integrating-user-pools-with-identity-pools.html>

upvoted 3 times

A company is building a scalable data management solution by using AWS services to improve the speed and agility of development. The solution will ingest large volumes of data from various sources and will process this data through multiple business rules and transformations.


The solution requires business rules to run in sequence and to handle reprocessing of data if errors occur when the business rules run. The company needs the solution to be scalable and to require the least possible maintenance.

Which AWS service should the company use to manage and automate the orchestration of the data flows to meet these requirements?

- A. AWS Batch
- B. AWS Step Functions
- C. AWS Glue
- D. AWS Lambda

Correct Answer: B

Community vote distribution



geekdamsel Highly Voted 1 year, 8 months ago

Got this question in exam. Correct answer is B.

upvoted 13 times

sumanshu Most Recent 1 month, 1 week ago

Selected Answer: B

AWS Step Functions is an orchestration service that is designed to manage and automate workflows

upvoted 1 times

sumanshu 1 month, 1 week ago

A) Eliminated - AWS Batch is great for running batch processing jobs, especially for large-scale computational workloads. However, it is not designed for orchestrating complex workflows

upvoted 1 times

sumanshu 1 month, 1 week ago

B) Eliminated - AWS Glue is a managed ETL (Extract, Transform, Load) service that helps you prepare and transform data. While it can handle data processing, it does not provide the orchestration and workflow management features

upvoted 1 times

sumanshu 1 month, 1 week ago

Sorry I mean C for above explanation

D) Eliminated - it does not provide built-in orchestration or workflow management capabilities like AWS Step Functions

upvoted 1 times

trieudo 1 month, 3 weeks ago

Selected Answer: B

I think those keyword when pick B: scalable, sequence, reprocessing of data if errors, least possible maintenance

upvoted 1 times

NagaoShingo 8 months ago

Selected Answer: B

B is correct answer.

upvoted 1 times

65703c1 8 months, 2 weeks ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

Dikshika 9 months, 3 weeks ago

Selected Answer: B

its clearly mention orchestration , sequence and multiple processing and transformations

upvoted 2 times

🗨️ **vinfo** 9 months, 4 weeks ago

Selected Answer: B

StepFunctions, es el servicio recomendado para orquestar. B correcto

upvoted 1 times

🗨️ **ibratoev** 10 months, 1 week ago

Selected Answer: B

B is correct

upvoted 1 times

🗨️ **alven_alinan** 1 year, 2 months ago

Selected Answer: B

Answer is B. Step Function is about orchestrating workflows

upvoted 3 times

🗨️ **dongocanh272** 1 year, 2 months ago

Selected Answer: B

My answer is B

upvoted 1 times

🗨️ **Digo30sp** 1 year, 3 months ago

Selected Answer: B

B is correct

upvoted 1 times

🗨️ **NinjaCloud** 1 year, 4 months ago

Best option: B

upvoted 1 times

🗨️ **panoptica** 1 year, 4 months ago

Selected Answer: B

b init

upvoted 1 times

🗨️ **sharma_ps93** 1 year, 4 months ago

The answer is B(Step Functions). For people confused with AWS Lambda, it is a compute service and can be used within Step Functions, but it alone does not provide the orchestration and error handling features required in this case.

upvoted 3 times

🗨️ **casharan** 1 year, 5 months ago

Selected Answer: D

check the link below:

<https://docs.aws.amazon.com/lambda/latest/operatorguide/orchestration.html>

upvoted 1 times

🗨️ **hmdev** 1 year, 5 months ago

Selected Answer: B

You can use Step functions to create a workflow of functions that should be invoked in a sequence. You can also push output from one one-step function and use it as an input for next-step function. Also, Step functions have very useful Retry and Catch -> error-handling features.

upvoted 1 times

🗨️ **jayvarma** 1 year, 5 months ago

Keyword: run in sequence and to handle reprocessing of data. So, answer is option B. And also each task in a step function can be handled by a different AWS Service such as AWS Lambda or AWS Glue which is used for ETL jobs.

upvoted 1 times

A developer has created an AWS Lambda function that is written in Python. The Lambda function reads data from objects in Amazon S3 and writes data to an Amazon DynamoDB table. The function is successfully invoked from an S3 event notification when an object is created. However, the function fails when it attempts to write to the DynamoDB table. What is the MOST likely cause of this issue?

- A. The Lambda function's concurrency limit has been exceeded.
- B. DynamoDB table requires a global secondary index (GSI) to support writes.
- C. The Lambda function does not have IAM permissions to write to DynamoDB.
- D. The DynamoDB table is not running in the same Availability Zone as the Lambda function.

Correct Answer: C

Community vote distribution

C (100%)

🗨️ **sumanshu** 1 month, 1 week ago

Selected Answer: C

A) Eliminated - If the function was invoked successfully from the S3 event notification but fails during the write operation, this is less likely to be the issue
upvoted 1 times

🗨️ **sumanshu** 1 month, 1 week ago

B) Eliminated - If the function was invoked successfully from the S3 event notification but fails during the write operation, this is less likely to be the issue
upvoted 1 times

🗨️ **sumanshu** 1 month, 1 week ago

D) Eliminated - AWS services, including Lambda and DynamoDB, are designed to work across Availability Zones without issues.
upvoted 1 times

🗨️ **trieudo** 1 month, 3 weeks ago

Selected Answer: C

keywords:

- fails when it attempts to write => discard A (when invoke success, not due to rate limit)

Knowledge:

- GSI relevant to read action, not write => discard B

- DynamoDB is global server, so AZ is not effect on it => discard D

upvoted 2 times

🗨️ **gomodhara** 1 month, 3 weeks ago

Selected Answer: C

Create an IAM policy that grants your Lambda function read and write permissions to a specific DynamoDB table by defining the actions dynamodb:GetItem, dynamodb:PutItem, and others as needed

upvoted 1 times

🗨️ **Aditya_bb_sharma** 3 months, 3 weeks ago

Selected Answer: C

Answer is B

upvoted 1 times

🗨️ **ACurryDeveloper** 6 months ago

Its c, who on earth would vote for D??

upvoted 1 times

🗨️ **nkroker** 7 months ago

C is correct because the question does not contain any details about the placement of the resources in different availability zones hence making it obvious for users to assume they are running in the same region and that's the reason why it makes sense that the Lambda execution role

does not have the write permission for dynamodb table.

upvoted 1 times

🗨️ **MessiVN** 7 months, 2 weeks ago

Selected Answer: C

C is correct

upvoted 1 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗨️ **Dikshika** 9 months, 3 weeks ago

Selected Answer: C

it mentions lambda is unable to write to Dynamodb and C seems most logical answer here

upvoted 1 times

🗨️ **vinfo** 9 months, 4 weeks ago

Selected Answer: C

C. es correcto. En general para estos comportamientos, se debe a temas de permisos.

upvoted 1 times

🗨️ **apa_1** 10 months, 1 week ago

Selected Answer: C

C is correct

upvoted 1 times

🗨️ **ibratoev** 10 months, 1 week ago

Selected Answer: C

It is C

upvoted 1 times

🗨️ **alven_alinan** 1 year, 2 months ago

Selected Answer: C

Answer is C

upvoted 1 times

🗨️ **dongocanh272** 1 year, 2 months ago

Why the correct answer is D? All of us think C must be the correct answer

upvoted 3 times

🗨️ **liddym2** 1 year, 2 months ago

Am i missing something? Why in God's name are the answer's provided wrong? It says D is the right answer. Its obviously C..

upvoted 3 times

🗨️ **dongocanh272** 1 year, 2 months ago

Selected Answer: C

I think C is correct.

upvoted 1 times

🗨️ **chvtjaswi** 1 year, 4 months ago

Selected Answer: C

correct answer is C

upvoted 3 times

A developer is creating an AWS CloudFormation template to deploy Amazon EC2 instances across multiple AWS accounts. The developer must choose the EC2 instances from a list of approved instance types.

How can the developer incorporate the list of approved instance types in the CloudFormation template?

- A. Create a separate CloudFormation template for each EC2 instance type in the list.
- B. In the Resources section of the CloudFormation template, create resources for each EC2 instance type in the list.
- C. In the CloudFormation template, create a separate parameter for each EC2 instance type in the list.
- D. In the CloudFormation template, create a parameter with the list of EC2 instance types as AllowedValues.

Correct Answer: D

Community vote distribution

D (100%)

🗳️ **Bibay** Highly Voted 1 year, 9 months ago

Selected Answer: D

Option D is the correct answer. In the CloudFormation template, the developer should create a parameter with the list of approved EC2 instance types as AllowedValues. This way, users can select the instance type they want to use when launching the CloudFormation stack, but only from the approved list.

Option A is not a scalable solution as it requires creating a separate CloudFormation template for each EC2 instance type, which can become cumbersome and difficult to manage as the number of approved instance types grows.

Option B is not necessary as creating resources for each EC2 instance type in the list would not enforce the requirement to choose only from the approved list. It would also increase the complexity of the template and make it difficult to manage.

Option C is not ideal as it would require creating a separate parameter for each EC2 instance type, which can become difficult to manage as the number of approved instance types grows. Also, it does not enforce the requirement to choose only from the approved list.

upvoted 24 times

🗳️ **jipark** 1 year, 6 months ago

quite much clear explanation !!!

upvoted 1 times

🗳️ **geekdamsel** Highly Voted 1 year, 8 months ago

Got this question in exam. Correct answer is D.

upvoted 8 times

🗳️ **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: D

A) Eliminated - higher maintenance overhead for maintaining multiple templates

B) Eliminated - This would create multiple EC2 instances unnecessarily, which does not align with the requirement to choose a single instance type from a list.

upvoted 1 times

🗳️ **trieudo** 1 month, 3 weeks ago

Selected Answer: D

==> Discard A: duplicate code, hard to maintain

==> Discard B: all resource will be created instead of it is necessary or not

==> Discard C: multiple param ==> when have larger param count, hard to maintain

D for dynamic for fixed parameter ==> most generic

upvoted 1 times

🗳️ **LocNV** 4 months, 1 week ago

Selected Answer: D

Parameters:

InstanceType:

Type: String

Default: 't2.micro'

AllowedValues:

- 't2.micro'

- 't2.small'

- 't2.medium'

- 't3.micro'

- 't3.small'

- 't3.medium'

Description: 'Select the EC2 instance type for deployment.'

Resources:

MyEC2Instance:

Type: 'AWS::EC2::Instance'

Properties:

ImageId: ami-12345678

InstanceType: !Ref InstanceType

upvoted 4 times

🗨️ **MessiVN** 7 months, 2 weeks ago

Selected Answer: D

D is correct

upvoted 1 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

🗨️ **vinfo** 9 months, 4 weeks ago

Selected Answer: D

Correcta D. Con este parámetro se permite dar permisos de elegir el tipo de instancia.

upvoted 1 times

🗨️ **apa_1** 10 months, 1 week ago

Selected Answer: D

option D is correct

upvoted 1 times

🗨️ **payireb682** 1 year, 1 month ago

Selected Answer: D

D is correct

upvoted 1 times

🗨️ **leonardoliveros** 1 year, 2 months ago

Selected Answer: D

D is the correct, because you are restricting the possible options to that parameter

upvoted 1 times

🗨️ **Pupina** 1 year, 7 months ago

Why B instead of C? Each AWS SDK implements retry logic automatically. Most AWS SDKs now support exponential backoff and jitter as part of their retry behavior

Then D to increase capacity <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/TroubleshootingThrottling.html>

C&D

upvoted 1 times

🗨️ **Pupina** 1 year, 7 months ago

This answer is for question 7 not 6

upvoted 1 times

🗨️ **NanaDanso** 1 year, 9 months ago

Selected Answer: D

D looks about right

upvoted 4 times

  **prabhay786** 1 year, 10 months ago

It should be D

upvoted 4 times

  **aragon_saa** 1 year, 10 months ago

D

<https://www.examttopics.com/discussions/amazon/view/88788-exam-aws-certified-developer-associate-topic-1-question-343/>


upvoted 3 times

A developer has an application that makes batch requests directly to Amazon DynamoDB by using the BatchGetItem low-level API operation. The responses frequently return values in the UnprocessedKeys element. Which actions should the developer take to increase the resiliency of the application when the batch response includes values in UnprocessedKeys? (Choose two.)

- A. Retry the batch operation immediately.
- B. Retry the batch operation with exponential backoff and randomized delay.
- C. Update the application to use an AWS software development kit (AWS SDK) to make the requests.
- D. Increase the provisioned read capacity of the DynamoDB tables that the operation accesses.
- E. Increase the provisioned write capacity of the DynamoDB tables that the operation accesses.

Correct Answer: BD

Community vote distribution



brandon87 Highly Voted 1 year, 10 months ago

Selected Answer: BD

(B) If you delay the batch operation using exponential backoff, the individual requests in the batch are much more likely to succeed.

(D) The most likely cause of a failed read or a failed write is throttling. For BatchGetItem, one or more of the tables in the batch request does not have enough provisioned read capacity to support the operation

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

upvoted 22 times

Untamables Highly Voted 1 year, 10 months ago

Selected Answer: BC

B & C

<https://docs.aws.amazon.com/general/latest/gr/api-retries.html>

upvoted 19 times

koniczny69 1 year ago

C already handles retries, why would want to do that manually?

upvoted 1 times

wtf3344 Most Recent 2 weeks, 2 days ago

Selected Answer: BD

BatchWriteItem

- UnprocessedItems for failed write operations (exponential backoff or add WCU)

BatchGetItem

- UnprocessedKeys for failed read operations (exponential backoff or add RCU)

So answer is B & D

upvoted 1 times

Hasitha99 3 weeks, 2 days ago

Selected Answer: BC

D also could be an answer. But is not the most suitable since there is a possibility of getting this result intermittently even though we have necessary capacity.

upvoted 1 times

Arad 1 month ago

Selected Answer: BC

BC is correct answer.


upvoted 1 times

sumanshu 1 month, 1 week ago

Selected Answer: BD

DynamoDB BatchGetItem API operation, the request fetches multiple items in one operation. However, DynamoDB has limits on the resources it allocates for processing requests. If it cannot process some items within a batch request due to these limits, it returns the processed items (the ones it successfully retrieved), and for the remaining items, it includes them in the UnprocessedKeys element in the response.

upvoted 2 times

 **trieudo** 1 month, 3 weeks ago

Selected Answer: BC

Keyword: increase, resiliency, application, response, UnprocessedKeys(Unprocessed due to be not enough infra)


==> Discard A: horrible way, push so many traffic to busy system

==> Discard D, E: Error due to huge traffic, it can be unlimited. Don't come from lacking resource ==> Your resource, your money is limit << unpredicted huge traffic

B: is good way to call x2 times after failed, call again after 1s -> 2s -> 4s -> 8s -> ... until success

C: is intergrate method of 'B' option automaically inside SDK

upvoted 1 times

 **9d8dd9c** 3 months, 2 weeks ago

Selected Answer: BD

BD

C handles retry but using SDK is not necessary here

upvoted 1 times

 **Phongsanth** 4 months, 1 week ago

Selected Answer: BD

I vote for B,D

Scroll down to the bottom of this page and you will see the reason. I paste some of words here.

The most likely cause of a failed read or a failed write is throttling. For BatchGetItem, one or more of the tables in the batch request does not have enough provisioned read capacity to support the operation. For BatchWriteItem, one or more of the tables does not have enough provisioned write capacity.

If DynamoDB returns any unprocessed items, you should retry the batch operation on those items. However, we strongly recommend that you use an exponential backoff algorithm. If you retry the batch operation immediately, the underlying read or write requests can still fail due to throttling on the individual tables. If you delay the batch operation using exponential backoff, the individual requests in the batch are much more likely to succeed.

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

upvoted 3 times

 **eberhe900** 4 months, 1 week ago

Selected Answer: BD

Backoff and Retry Strategy:

Implement a backoff and retry strategy to prevent overwhelming the DynamoDB service with repeated requests.

Apply an exponential backoff algorithm, where you progressively increase the delay between each retry attempt.

Consider implementing a maximum number of retries to avoid an infinite retry loop.

Fine-Tuning DynamoDB Provisioned Capacity:

If you consistently encounter unprocessed items during batch operations, it may indicate that your DynamoDB table's provisioned capacity is insufficient.

Monitor the table's consumed capacity and adjust the provisioned capacity (read capacity units) accordingly to handle the load and reduce the occurrence of unprocessed items.

By implementing these steps, you can effectively handle unprocessed items returned by the BatchGetItem operation in DynamoDB and ensure that all items are processed successfully.

upvoted 2 times

 **SherzodBek** 4 months, 1 week ago

Selected Answer: BD

B & D.

B is correct. Because in the question, it is mentioned that low-level API is being used. It means exponential backoff can be implemented manually.

D is correct. Because there is a frequently keyword in the question. If UnprocessedKeys error occurs frequently, DynamoDB doesn't have enough capacity to process requests. So read capacity should be increased.

upvoted 3 times

  **TheFivePips** 4 months, 1 week ago

Selected Answer: BD

Exponential backoff with randomized delay is a common technique used to handle transient failures and throttle errors in distributed systems like DynamoDB. This approach involves retrying the failed operation after waiting for an increasing amount of time, which helps reduce the load on the service and increases the likelihood of success during periods of high demand or throttling.

If the BatchGetItem operation frequently returns values in the UnprocessedKeys element, it indicates that the table's read capacity might be insufficient to handle the requested workload. By increasing the provisioned read capacity for the DynamoDB tables, the application can better handle the read throughput requirements and reduce the likelihood of encountering UnprocessedKeys in batch responses.



AWS SDK might provide additional features and simplifications for making requests, it does not directly address the issue of UnprocessedKeys in batch responses. This option might be beneficial for improving code maintainability and leveraging SDK features however.

upvoted 1 times

  **phongnx8** 6 months ago

BC is the correct answer.

upvoted 1 times



  **Anandesh** 6 months, 2 weeks ago

Selected Answer: BC

https://docs.aws.amazon.com/amazondynamodb/latest/APIReference/API_BatchGetItem.html

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



upvoted 1 times

  **65703c1** 8 months, 2 weeks ago

Selected Answer: BC

BC is the correct answer.

upvoted 1 times

  **drycleansing** 9 months, 3 weeks ago

Selected Answer: BC

among B C D, it is hard to say D copes with the problem directly I guess. Increasing RCU will affect the ratio of unprocessed items but that does not mean it handles the unprocessed items.

upvoted 1 times

  **vinfo** 9 months, 4 weeks ago

Selected Answer: BD

B,D. La combinación de estrategias es ideal para este comportamiento

upvoted 1 times

A company is running a custom application on a set of on-premises Linux servers that are accessed using Amazon API Gateway. AWS X-Ray tracing has been enabled on the API test stage.


How can a developer enable X-Ray tracing on the on-premises servers with the LEAST amount of configuration?

- A. Install and run the X-Ray SDK on the on-premises servers to capture and relay the data to the X-Ray service.
- B. Install and run the X-Ray daemon on the on-premises servers to capture and relay the data to the X-Ray service.
- C. Capture incoming requests on-premises and configure an AWS Lambda function to pull, process, and relay relevant data to X-Ray using the PutTraceSegments API call.
- D. Capture incoming requests on-premises and configure an AWS Lambda function to pull, process, and relay relevant data to X-Ray using the PutTelemetryRecords API call.

Correct Answer: B

Community vote distribution

B (100%)


 **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: B

B

<https://docs.aws.amazon.com/xray/latest/devguide/xray-daemon.html>

upvoted 8 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B

A) Eliminated - The X-Ray SDK is used to instrument your application code to create and send trace data to the X-Ray daemon.

However, this option alone is insufficient because the SDK only captures traces—it still needs the X-Ray daemon to relay data to the X-Ray service.

upvoted 2 times

 **sumanshu** 1 month, 1 week ago

C) - Eliminated - While this could work, but involves more configuration and complexity

upvoted 1 times

 **sumanshu** 1 month, 1 week ago

D) Eliminated - This API is not suitable for sending trace data.

upvoted 1 times

 **aloksheth** 1 month, 2 weeks ago

Selected Answer: B

B

<https://docs.aws.amazon.com/xray/latest/devguide/xray-daemon.html>

upvoted 1 times

 **trieudo** 1 month, 3 weeks ago

Selected Answer: B


Keyword: LEAST amount of configuration; accessed using Amazon API Gateway. AWS X-Ray tracing has been enabled

==> Discard C, D: must write logic in lambda funct extra

==> Discard A: must change code to add annotaion, If not have keyword: 'accessed using Amazon API Gateway. AWS X-Ray tracing has been enabled', it will be correct, because it easier than B

B: despite you must know deeply X-Ray to using it (config JSON format, ...), but with conditon in keyword, you just run only cmd './xray --region us-east-1' for working

upvoted 1 times

 **Hendrix944** 3 months, 3 weeks ago

Selected Answer: B

B is correct according to Stephane maarek course

upvoted 1 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗨️ **vinfo** 9 months, 4 weeks ago

Selected Answer: B

B. El daemon, es una capacidad propia de X-Ray, para instalar directamente

upvoted 1 times

🗨️ **apa_1** 10 months, 1 week ago

Selected Answer: B

B is correct

upvoted 1 times

🗨️ **leonardoliveros** 1 year, 2 months ago

Selected Answer: B

B, you should to install the X-Ray daemon in on-premises without this all others option is wrong

upvoted 1 times

🗨️ **Ugo_22** 1 year, 3 months ago

Selected Answer: B

The answer is obviously B.

upvoted 1 times

🗨️ **Kowalsky95** 1 year, 4 months ago

From doc: The AWS X-Ray daemon is a software application that listens for traffic on UDP port 2000, gathers raw segment data, and relays it to the AWS X-Ray API. The daemon works in conjunction with the AWS X-Ray SDKs and must be running so that data sent by the SDKs can reach the X-Ray service.

Running just the daemon won't achieve anything.

upvoted 3 times

🗨️ **geekdamsel** 1 year, 8 months ago

Got this question in exam. Correct answer is B.

upvoted 3 times

🗨️ **Bibay** 1 year, 9 months ago

Selected Answer: B

. Install and run the X-Ray daemon on the on-premises servers to capture and relay the data to the X-Ray service is the correct option. The X-Ray daemon can be installed and configured on the on-premises servers to capture data and send it to the X-Ray service. This requires minimal configuration and setup. Option A is incorrect because while the X-Ray SDK can be used to capture data on the on-premises servers, it requires more configuration and development effort than the X-Ray daemon. Option C and D are also incorrect because they involve setting up an AWS Lambda function, which is not necessary for enabling X-Ray tracing on the on-premises servers.

upvoted 2 times

🗨️ **ihta_2031** 1 year, 10 months ago

Selected Answer: B

It's B

upvoted 4 times

🗨️ **haaris786** 1 year, 10 months ago

B: It is Daemon which can be installed for Linux

upvoted 3 times

🗨️ **aragon_saa** 1 year, 10 months ago

B

<https://www.examttopics.com/discussions/amazon/view/28998-exam-aws-certified-developer-associate-topic-1-question-324/>

upvoted 3 times

A company wants to share information with a third party. The third party has an HTTP API endpoint that the company can use to share the information. The company has the required API key to access the HTTP API.

The company needs a way to manage the API key by using code. The integration of the API key with the application code cannot affect application performance.

Which solution will meet these requirements MOST securely?

- A. Store the API credentials in AWS Secrets Manager. Retrieve the API credentials at runtime by using the AWS SDK. Use the credentials to make the API call.
- B. Store the API credentials in a local code variable. Push the code to a secure Git repository. Use the local code variable at runtime to make the API call.
- C. Store the API credentials as an object in a private Amazon S3 bucket. Restrict access to the S3 object by using IAM policies. Retrieve the API credentials at runtime by using the AWS SDK. Use the credentials to make the API call.
- D. Store the API credentials in an Amazon DynamoDB table. Restrict access to the table by using resource-based policies. Retrieve the API credentials at runtime by using the AWS SDK. Use the credentials to make the API call.

Correct Answer: A

Community vote distribution


A (100%)

 **Kristijan92** Highly Voted 1 year, 10 months ago

Selected Answer: A

answer A

upvoted 13 times

 **elfinka9** Highly Voted 1 year, 6 months ago

Selected Answer: A


Why B is marked as correct ????

upvoted 9 times

 **ACurryDeveloper** 6 months ago

Why is developer being a benchod? A, b, c or d?

upvoted 1 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: A

AWS Secrets Manager is purpose-built for securely managing sensitive data such as API keys

upvoted 2 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: A

keyword: MOST securely

==> discard B first: If the code repository is compromised, the API key could be exposed.

==> discard C, D after: this requires setting up and managing the DynamoDB table and resource-based policies ==> can make mistake from ppl, then creating issue about security

A: AWS Secrets Manager is a specialized service designed to securely store sensitive information like API keys.

upvoted 2 times

 **Bibay** 4 months, 1 week ago

Selected Answer: A

The MOST secure solution to manage the API key while ensuring that the integration of the API key with the application code does not affect application performance is to store the API key in AWS Secrets Manager. The API key can be retrieved at runtime by using the AWS SDK, which does not impact application performance. Therefore, option A is the correct answer.

Option B is not secure as it exposes the API key to anyone with access to the code repository, which increases the risk of unauthorized access.



Option C and D may work, but they require additional configuration and permissions management. Storing the API key in an S3 bucket or a DynamoDB table could expose the key to unauthorized users if proper IAM policies are not in place. Therefore, option A is the most secure and simple solution to manage the API key while not affecting the application's performance.

upvoted 2 times

  **nkroker** 7 months ago

B is certainly a wrong answer because if you read the application development with AWS case study and also the best practices then in it the AWS itself does not encourage the developer to store the sensitive API and license keys in code on the other hand it encourages to use the provided solutions like KMS, Secrets Manager and Parameter Store.



upvoted 1 times

  **297dfbb** 7 months, 3 weeks ago

Selected Answer: A

A is correct. B is wrong because you never store credentials in source control.



upvoted 1 times

  **c949c84** 7 months, 3 weeks ago

Selected Answer: A

Most secure



upvoted 2 times

  **65703c1** 8 months, 2 weeks ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

  **Prosen2522** 8 months, 4 weeks ago

A is it

upvoted 1 times

  **badsati** 10 months ago

Straight A

upvoted 1 times

  **bednark** 1 year ago

Selected Answer: A

A is correct

upvoted 1 times



  **leonardoliveros** 1 year, 2 months ago

Selected Answer: A

B isn't security

A is the best option for this scenary



upvoted 2 times

  **gullyboy77** 1 year, 4 months ago

Selected Answer: A

Secret Manager is the safest way to store secrets in AWS.



upvoted 2 times

  **chvtjaswi** 1 year, 4 months ago

Selected Answer: A

Answer A

upvoted 2 times

  **hmdev** 1 year, 5 months ago

Selected Answer: A

A seems to be the most secure and correct. Always use Secret Manger to store secrets, as the name implies.

upvoted 1 times

  **Yuxing_Li** 1 year, 5 months ago

Selected Answer: A

A is correct

upvoted 1 times

A developer is deploying a new application to Amazon Elastic Container Service (Amazon ECS). The developer needs to securely store and retrieve different types of variables. These variables include authentication information for a remote API, the URL for the API, and credentials. The authentication information and API URL must be available to all current and future deployed versions of the application across development, testing, and production environments.

How should the developer retrieve the variables with the FEWEST application changes?

- A. Update the application to retrieve the variables from AWS Systems Manager Parameter Store. Use unique paths in Parameter Store for each variable in each environment. Store the credentials in AWS Secrets Manager in each environment.
- B. Update the application to retrieve the variables from AWS Key Management Service (AWS KMS). Store the API URL and credentials as unique keys for each environment.
- C. Update the application to retrieve the variables from an encrypted file that is stored with the application. Store the API URL and credentials in unique files for each environment.
- D. Update the application to retrieve the variables from each of the deployed environments. Define the authentication information and API URL in the ECS task definition as unique names during the deployment process.


Correct Answer: A

Community vote distribution

A (100%)


 **geekdamsel** Highly Voted 1 year, 8 months ago

Correct answer is A.
upvoted 18 times

 **Warlord_92** Highly Voted 1 year, 10 months ago

Selected Answer: A

The application has credentials and URL, so it's convenient to store them in ssm parameter store retrieve them.
upvoted 9 times


 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: A

B) Eliminated - It cannot directly store variables such as the API URL or credentials
D) Eliminated - Storing sensitive information like credentials directly in ECS task definitions is not secure
upvoted 1 times

 **sumanshu** 1 month, 1 week ago


C) Eliminated - Storing variables in encrypted files adds operational overhead. Managing separate files for each environment can quickly become cumbersome.
upvoted 1 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: A

keyword: FEWEST application changes

=> A: This minimizes code changes as the retrieval method remains consistent across environments; only the Parameter Store paths need updating. Secrets Manager securely stores sensitive credentials.
upvoted 1 times

 **Tee400** 4 months, 1 week ago

Selected Answer: A

AWS Systems Manager Parameter Store is a service that allows you to securely store configuration data such as API URLs, credentials, and other variables. By updating the application to retrieve the variables from Parameter Store, you can separate the configuration from the application code, making it easier to manage and update the variables without modifying the application itself. Storing the credentials in AWS Secrets Manager provides an additional layer of security for sensitive information.
upvoted 2 times

 **badsati** 4 months, 1 week ago

Answer is A

AWS Systems Manager Parameter Store and AWS Secrets Manager are designed for securely storing and managing sensitive information such as credentials, API URLs, and configuration variables.

Using AWS Systems Manager Parameter Store allows the developer to centrally manage configuration variables across different environments (development, testing, production) without requiring application code changes. Each variable can have a unique path in Parameter Store, ensuring separation and organization.


upvoted 2 times

 **ez_24** 4 months, 1 week ago

Correct Answer is A

Option B, using AWS Key Management Service (AWS KMS), is not ideal for this scenario primarily because AWS KMS is designed for creating and controlling encryption keys, not for storing configuration data or credentials. KMS keys are used to encrypt and decrypt data, rather than directly storing or managing it. For securely managing and retrieving application configuration data and sensitive information like API credentials, Systems Manager Parameter Store and AWS Secrets Manager are more appropriate, offering direct support for these use cases with better integration for applications.


upvoted 5 times

 **65703c1** 8 months, 2 weeks ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

 **Prosen2522** 8 months, 4 weeks ago

Selected Answer: A

SSM parameter store with proper posh is the answer.

upvoted 1 times

 **leonardoliveros** 1 year, 2 months ago

Selected Answer: A


You put the different variables for each environment, is the best solution because it's isolated between environment

upvoted 1 times

 **vmintam** 1 year, 3 months ago

i think corrent is A, but why is B ?

upvoted 2 times

 **alihaidar907** 1 year, 4 months ago

I think the wording of option A has a typo first it mentioned " Update the application to retrieve the variables from AWS Systems Manager Parameter Store" then it says "Store the credentials in AWS Secrets Manager in each environment."

upvoted 1 times

 **meetparag81** 1 year, 5 months ago


A is correct

upvoted 1 times

 **jayvarma** 1 year, 5 months ago

Option A is correct. The AWS Systems Manager Paramter Store's primary purpose is to secure sensitive information such as API URLs, credentials and the variables that we store in it.


upvoted 2 times

 **MrTee** 1 year, 9 months ago

Selected Answer: A


his solution allows the developer to securely store and retrieve different types of variables, including authentication information for a remote API, the URL for the API, and credentials.

upvoted 2 times

 **qsergii** 1 year, 9 months ago

Definitely A

upvoted 1 times

 **fqmark** 1 year, 9 months ago

it should be a, kms is used for encryption: <https://aws.amazon.com/kms/>

upvoted 3 times

A company is migrating legacy internal applications to AWS. Leadership wants to rewrite the internal employee directory to use native AWS services. A developer needs to create a solution for storing employee contact details and high-resolution photos for use with the new application.

Which solution will enable the search and retrieval of each employee's individual details and high-resolution photos using AWS APIs?

- A. Encode each employee's contact information and photos using Base64. Store the information in an Amazon DynamoDB table using a sort key.
- B. Store each employee's contact information in an Amazon DynamoDB table along with the object keys for the photos stored in Amazon S3.
- C. Use Amazon Cognito user pools to implement the employee directory in a fully managed software-as-a-service (SaaS) method.
- D. Store employee contact information in an Amazon RDS DB instance with the photos stored in Amazon Elastic File System (Amazon EFS).

Correct Answer: B

Community vote distribution

B (100%)


 **Bibay** Highly Voted 1 year, 9 months ago

Selected Answer: B

B. Store each employee's contact information in an Amazon DynamoDB table along with the object keys for the photos stored in Amazon S3.

Storing each employee's contact information in an Amazon DynamoDB table along with the object keys for the photos stored in Amazon S3 provides a scalable and efficient solution for storing and retrieving employee details and high-resolution photos using AWS APIs. The developer can use the DynamoDB table to query and retrieve employee details, while the S3 bucket can be used to store the high-resolution photos. By using S3, the solution can support large amounts of data while enabling fast retrieval times. The combination of DynamoDB and S3 can provide a cost-effective and scalable solution for storing employee data and photos.

upvoted 9 times


 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B

A) Eliminated - DynamoDB is not cost-effective for storing large binary data and also increases query latency

C) Eliminated - Cognito user pools are intended for user authentication and authorization, not as a storage solution

upvoted 1 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: B

==> discard A: store image in db, will enhance time query

==> discard C: cognito for identifying not querying info

==> discard D: EFS vs S3, I choose s3, its storage is global, best popular to store image than EFS, with EFS you must use EC2, how you can link image in EC2 with db

==> B is best choice

upvoted 1 times

 **ACurryDeveloper** 6 months ago

A: Base 64 is a distraction. You can use encryption at rest using KMS for most things. Would you store photos in DynamoDB? Would be silly

C. Cognito has nothing to do with the question. Question is asking about searching for employee details

D. Could work, but seems convoluted. Relational DBs are easily searchable, but how would you link the db to the image in EC2?

B is the simplest and correct answer here.


upvoted 1 times

 **65703c1** 8 months, 2 weeks ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

  **badsati** 9 months, 3 weeks ago

Selected Answer: B

Answer is B

upvoted 1 times

  **Baalhammun** 11 months, 3 weeks ago

Selected Answer: B

I agree, B is correct, DynamoDB to store user's data along the Key for S3 objects knowing that S3 is a good solution to store large amount of data or "high quality" images



upvoted 1 times

  **leonardoliveros** 1 year, 2 months ago

Selected Answer: B

DynamoDb + S3 is the best option for those scenarios

upvoted 1 times

  **hmdev** 1 year, 5 months ago

Selected Answer: B

DynamoDB is very fast, secure, and scalable. The S3 is very in-expensive, virtually limitless, and can handle large files. So B is the correct answer.

upvoted 2 times

  **ninomfr64** 1 year, 5 months ago

Selected Answer: B

A. is not really clear to me, however encoding all info in base64 would make search a bit complex

C. does not provide a solution for high resolution image

D. EFS does not provide API access to content

upvoted 3 times

  **jayvarma** 1 year, 5 months ago

Option B. As the question says that we have to store high-resolution photos, the solution is to use the S3 here. Because, DynamoDb cannot be used to store anything that is above 400 KB for each object.

In this case, we can use DynamoDb to store the contact information of each of the employees and reference the object keys in the table to retrieve the high-resolution images.

upvoted 2 times

  **ihta_2031** 1 year, 10 months ago

Selected Answer: B

Agreed with B

upvoted 4 times

  **aragon_saa** 1 year, 10 months ago

B

<https://www.examttopics.com/discussions/amazon/view/88823-exam-aws-certified-developer-associate-topic-1-question-240/>

upvoted 4 times

A developer is creating an application that will give users the ability to store photos from their cellphones in the cloud. The application needs to support tens of thousands of users. The application uses an Amazon API Gateway REST API that is integrated with AWS Lambda functions to process the photos. The application stores details about the photos in Amazon DynamoDB.

Users need to create an account to access the application. In the application, users must be able to upload photos and retrieve previously uploaded photos. The photos will range in size from 300 KB to 5 MB.


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

- A. Use Amazon Cognito user pools to manage user accounts. Create an Amazon Cognito user pool authorizer in API Gateway to control access to the API. Use the Lambda function to store the photos and details in the DynamoDB table. Retrieve previously uploaded photos directly from the DynamoDB table.
- B. Use Amazon Cognito user pools to manage user accounts. Create an Amazon Cognito user pool authorizer in API Gateway to control access to the API. Use the Lambda function to store the photos in Amazon S3. Store the object's S3 key as part of the photo details in the DynamoDB table. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.
- C. Create an IAM user for each user of the application during the sign-up process. Use IAM authentication to access the API Gateway API. Use the Lambda function to store the photos in Amazon S3. Store the object's S3 key as part of the photo details in the DynamoDB table. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.
- D. Create a users table in DynamoDB. Use the table to manage user accounts. Create a Lambda authorizer that validates user credentials against the users table. Integrate the Lambda authorizer with API Gateway to control access to the API. Use the Lambda function to store the photos in Amazon S3. Store the object's S3 key as part of the photo details in the DynamoDB table. Retrieve previously uploaded photos by querying DynamoDB for the S3 key.

Correct Answer: B

Community vote distribution

B (100%)

 **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: B

B

<https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-integrate-with-cognito.html>

<https://aws.amazon.com/blogs/big-data/building-and-maintaining-an-amazon-s3-metadata-index-without-servers/>

upvoted 12 times


 **jayvarma** Highly Voted 1 year, 5 months ago

As it is not a good practice to create a new IAM user for each user that signs up for the application, Option C is ruled out. Amazon Cognito user pools primary purpose is to authenticate and authorize web and mobile applications.

As the solution requires the application to store images that are between 300KB and 5MB in size, The idea of storing the images in the DynamoDB is ruled out because the object size in a dynamoDb table cannot exceed 400kb. The ideal solution for this problem would be to store the photos in S3 and store the object's key in the DynamoDB table.

So, Option B is the right answer

upvoted 6 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B

A) Eliminated - DynamoDB is not optimized for large binary objects

C) Eliminated - Creating IAM users for each application user is not scalable

D) Eliminated - Using a custom table and Lambda authorizer for user authentication increases operational complexity.

upvoted 1 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: B

keyword: LEAST operational overhead, size from 300 KB to 5 MB.

=> discard A: item limit upto 400kB only

==> discard C: each user per IAM user --> really bad practice

==> discard D: violate "LEAST operational overhead", you must build a lot (write lambda function), result in manage it alot

B is best choice to use services: s3, dynamoDB, Amazon Cognito with maximum automatic build-in feature and least effort for operation
upvoted 1 times

🗨️ **Saudis** 3 months ago

The difference between B and C the key word create user > cognito

upvoted 1 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗨️ **badsati** 9 months, 3 weeks ago

Selected Answer: B

No question, Answer is B

upvoted 1 times

🗨️ **TheFivePips** 11 months ago

Selected Answer: B

Definitely do not add all users manually, so that rules out C and D.

You wouldnt use DynamoDB to store the photos because DynamoDB limits the size of each item that you store in a table to 400 KB. So that rules out A

upvoted 2 times

🗨️ **leonardoliveros** 1 year, 2 months ago

Selected Answer: B

It's easier if you leverage all pros of Amazon Cognito you don't need creating a IAM user by employee

upvoted 1 times

🗨️ **ihta_2031** 1 year, 10 months ago

Selected Answer: B

Cognito,

Item size in dynamodb is less than this scenario

upvoted 4 times

🗨️ **pratchatcap** 1 year, 10 months ago

Selected Answer: B

B is the most valid solution.

A nearest, but invalid, because you cannot store object in Dynamo.

upvoted 3 times

A company receives food orders from multiple partners. The company has a microservices application that uses Amazon API Gateway APIs with AWS Lambda integration. Each partner sends orders by calling a customized API that is exposed through API Gateway. The API call invokes a shared Lambda function to process the orders.

Partners need to be notified after the Lambda function processes the orders. Each partner must receive updates for only the partner's own orders. The company wants to add new partners in the future with the fewest code changes possible.

Which solution will meet these requirements in the MOST scalable way?


- A. Create a different Amazon Simple Notification Service (Amazon SNS) topic for each partner. Configure the Lambda function to publish messages for each partner to the partner's SNS topic.
- B. Create a different Lambda function for each partner. Configure the Lambda function to notify each partner's service endpoint directly.
- C. Create an Amazon Simple Notification Service (Amazon SNS) topic. Configure the Lambda function to publish messages with specific attributes to the SNS topic. Subscribe each partner to the SNS topic. Apply the appropriate filter policy to the topic subscriptions.
- D. Create one Amazon Simple Notification Service (Amazon SNS) topic. Subscribe all partners to the SNS topic.

Correct Answer: C

Community vote distribution

C (78%)

A (22%)

 **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: C

C

<https://docs.aws.amazon.com/sns/latest/dg/sns-message-filtering.html>


upvoted 11 times

 **Bibay** Highly Voted 4 months, 1 week ago

Selected Answer: C

Option C is the most scalable way to meet the requirements. This solution allows for a single SNS topic to be used for all partners, which minimizes the need for code changes when adding new partners. By publishing messages with specific attributes to the SNS topic and applying the appropriate filter policy to the topic subscriptions, partners will only receive notifications for their own orders. This approach allows for a more flexible and scalable solution, where new partners can be added to the system with minimal changes to the existing codebase. Option A and D may not be scalable when there are a large number of partners, as creating a separate SNS topic for each partner or subscribing all partners to a single topic may not be feasible. Option B may result in a large number of Lambda functions that need to be managed separately.

upvoted 7 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: C

A) Eliminated - Every time a new partner is added, a new SNS topic needs to be created. This requires manual configuration and updates to the infrastructure.

upvoted 1 times

 **sumanshu** 1 month, 1 week ago

B) Eliminated - Adding a new partner means creating a new Lambda function, increasing operational overhead.

upvoted 1 times

 **sumanshu** 1 month, 1 week ago

C) Correct - A single SNS topic is used. The Lambda function publishes messages with attributes (e.g., partner ID). Each partner subscribes to the SNS topic and uses a filter policy to only receive messages relevant to their partner ID

upvoted 1 times

 **sumanshu** 1 month, 1 week ago

D) Eliminated - All partners would receive updates for all orders because there is no filtering mechanism to ensure partner-specific messages.

upvoted 1 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: C

keyword: MOST scalable way, only the partner's own orders



==> Discard A: you must update lambda for new topic added, but it simple for case having few partners, and little change

==> Discard B: You must create duplicate lambda function and maintain it. But best case for customized complexity requirements

==> Discard D: it is easy to scalable, but violate rule 'only the partner's own orders', when a partner can see msg of all others

C: best choice, match with 2 keywords above

upvoted 1 times

  **wail1997** 5 months, 1 week ago

Selected Answer: A

because of the fact that By default, you can have up to 200 filter policies per topic, the C option can't be the wright answer, but it's the A choice. since we can go up to 100 00 topics per SNS

upvoted 1 times

  **AnthonyTL** 4 months, 3 weeks ago



With C, if it's more that 200 partners, we could create another SNS for the next 200 partners. so it couls support up to 2000000 partners.

upvoted 2 times

  **ACurryDeveloper** 6 months ago

A works, but C is better, benchods. More efficient. Have to remember good curry cannot be had in a hurry



upvoted 1 times

  **Anandesh** 6 months, 2 weeks ago

Selected Answer: C

<https://docs.aws.amazon.com/sns/latest/dg/example-filter-policies.html>



upvoted 1 times

  **65703c1** 8 months, 2 weeks ago

Selected Answer: C

C is the correct answer.



upvoted 1 times

  **Prosen2522** 8 months, 4 weeks ago

Selected Answer: C

Funny understand why some people want to create separate SNS for each partner. You have got the option to filter and send notifications to the appropriate partner.

upvoted 1 times

  **badsati** 9 months, 3 weeks ago

Selected Answer: C

Answer is C ... No Question

upvoted 1 times

  **xdkonorek2** 1 year, 1 month ago

Selected Answer: A

you can create up to



10.000 filter policies per AWS account

200 filter policies per topic (not subscription!) limits option C to 200 partners

100 000 topics per AWS account, limits option A to 100 000 partners

A and C works but A has better scalability with ability to add 100 000 partners

upvoted 5 times

  **drycleansing** 9 months, 3 weeks ago

the best answer

upvoted 1 times

  **leonardoliveros** 1 year, 2 months ago

Selected Answer: C

You can using a filter policy to just sent the info by partner

upvoted 1 times

  **ninomfr64** 1 year, 5 months ago

Selected Answer: C

C. adding a new partner would only require to create a new subscription with the right filter

upvoted 1 times

🗨️ **tttamttam** 1 year, 6 months ago

Selected Answer: C

C seems the most efficient way. when you add more partners, you can just assign new codes for each partner. with the codes, you can send notifications to specific paters

upvoted 1 times

🗨️ **rld2000** 1 year, 6 months ago

Selected Answer: A

The answer is A since this question has two crucial requirements:

a) ... with the fewest code changes possible.

b) ...in the MOST scalable way

ChatGPT initially gives an incorrect answer and then adjusts its response when requirements are asked.

upvoted 1 times

🗨️ **rld2000** 1 year, 6 months ago

OOH another important requirement: Each partner must receive updates for only the partner's own orders, that is not achievable with option C

upvoted 1 times

🗨️ **Jeremy11** 1 year, 6 months ago

This part of C seems to meet that requirement: Apply the appropriate filter policy to the topic subscriptions.

upvoted 5 times

🗨️ **Skywalker23** 1 year, 4 months ago

Cannot be A. It requires change of lambda function code to send notifications to new SNS topics for new partners. Not a scalable solution.

upvoted 2 times

🗨️ **geekdamsel** 1 year, 8 months ago

Got this question in exam. Correct answer is C.

upvoted 4 times

🗨️ **Rpod** 1 year, 9 months ago

Selected Answer: C

C is the answer

upvoted 2 times

A financial company must store original customer records for 10 years for legal reasons. A complete record contains personally identifiable information (PII). According to local regulations, PII is available to only certain people in the company and must not be shared with third parties. The company needs to make the records available to third-party organizations for statistical analysis without sharing the PII. A developer wants to store the original immutable record in Amazon S3. Depending on who accesses the S3 document, the document should be returned as is or with all the PII removed. The developer has written an AWS Lambda function to remove the PII from the document. The function is named `removePii`.

What should the developer do so that the company can meet the PII requirements while maintaining only one copy of the document?

- A. Set up an S3 event notification that invokes the `removePii` function when an S3 GET request is made. Call Amazon S3 by using a GET request to access the object without PII.
- B. Set up an S3 event notification that invokes the `removePii` function when an S3 PUT request is made. Call Amazon S3 by using a PUT request to access the object without PII.
- C. Create an S3 Object Lambda access point from the S3 console. Select the `removePii` function. Use S3 Access Points to access the object without PII.
- D. Create an S3 access point from the S3 console. Use the access point name to call the `GetObjectLegalHold` S3 API function. Pass in the `removePii` function name to access the object without PII.

Correct Answer: C

Community vote distribution

C (100%)

 **gcmrjbr** Highly Voted 1 year, 2 months ago


An S3 Object Lambda access point is a new type of access point that you can create to invoke your own AWS Lambda function to modify the content of an S3 object. You can use S3 Object Lambda access points to transform data as it is being retrieved from an S3 bucket, without modifying the original data stored in the bucket

upvoted 15 times

 **Skip** 7 months, 1 week ago

Thanks for the info. Great heads up!

upvoted 1 times


 **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: C

C

<https://aws.amazon.com/s3/features/object-lambda/>

upvoted 12 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: C

A) Eliminated - function cannot be invoked when a GET request is made.

B) Eliminated - It will either create two copies or overwrite existing copy

D) Eliminated - `GetObjectLegalHold` this API is used to check if an object is under a legal hold. It has nothing to do with dynamically modifying or removing PII from documents.

upvoted 2 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: C

=> Discard A: S3 event notification not support GET

=> Discard B: violate this rule 'keyword: only one copy of the document'. S3 event notification with PUT, make a copy (without PII) beside original record

=> Discard D: API `GetObjectLegalHold` is s3 GET API to know s3 object is editable, it can't update/ edit s3 object. Lambda can't auto be called by this API

C: Lambda access point is intermediary s3 object and end-user, it modifies a copy of data (delete PII) then return user, then delete this data from local lambda memory

upvoted 1 times

🗨️ 👤 **ahadh7621** 6 months, 2 weeks ago

Selected Answer: C

<https://aws.amazon.com/s3/features/object-lambda/>

With S3 Object Lambda, you can add your own code to S3 GET, HEAD, and LIST requests to modify and process data as it is returned to an application. You can use custom code to modify the data returned by S3 GET requests to filter rows, dynamically resize images, redact confidential data, and much more.

upvoted 2 times

🗨️ 👤 **Anandesh** 6 months, 2 weeks ago

Selected Answer: C

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/transforming-objects.html>

upvoted 1 times

🗨️ 👤 **65703c1** 8 months, 2 weeks ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗨️ 👤 **KarBiswa** 10 months, 3 weeks ago

Selected Answer: C

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/olap-create.html>

upvoted 1 times

🗨️ 👤 **pagyabeng** 1 year, 8 months ago

Why is it C?

upvoted 3 times

🗨️ 👤 **geekdamsel** 1 year, 8 months ago

Correct answer is C.

upvoted 2 times

🗨️ 👤 **Rpod** 1 year, 9 months ago

Selected Answer: C

C answer

upvoted 1 times

🗨️ 👤 **ihta_2031** 1 year, 10 months ago

Selected Answer: C

It is C

upvoted 3 times

🗨️ 👤 **aragon_saa** 1 year, 10 months ago

C

<https://www.examttopics.com/discussions/amazon/view/88229-exam-aws-certified-developer-associate-topic-1-question-174/>

upvoted 7 times

A developer is deploying an AWS Lambda function. The developer wants the ability to return to older versions of the function quickly and seamlessly.

How can the developer achieve this goal with the LEAST operational overhead?

- A. Use AWS OpsWorks to perform blue/green deployments.
- B. Use a function alias with different versions.
- C. Maintain deployment packages for older versions in Amazon S3.
- D. Use AWS CodePipeline for deployments and rollbacks.

Correct Answer: B

Community vote distribution

B (100%)

🗨️ **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: B

B

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html>

upvoted 7 times

🗨️ **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B

A) Eliminated - OpsWorks is unnecessary for Lambda function deployments. Blue/green deployments are more suited for EC2 or ECS applications, not for a simple Lambda rollback.

B) Correct - Quickly switch the alias to an older version

C) Eliminated - Requires manual work to download and redeploy older versions.

D) Eliminated - CodePipeline introduces unnecessary complexity

upvoted 1 times

🗨️ **trieudo** 1 month, 2 weeks ago

Selected Answer: B

keyword: LEAST operational overhead

=> discard D: it does more work, so cost for many works also increase

=> discard C: it makes effort for manual work, adding operational overhead and delays.

=> Discard A: Designed for managing server configurations, not for managing lambda function

B: most seamless, only change version by pointing an alias (eg: production) to specific Lambda Version. On the other hand, it's free

upvoted 1 times

🗨️ **f271c23** 1 month, 3 weeks ago

Selected Answer: A

I am not sure if the question is about which deployment strategy to choose from if we want to roll back seamlessly once deployed. Or is the question about how to manage the deployment versions to be able to roll back. does anyone has similar doubts ? please help , thanks

upvoted 1 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: B

B is the correct answer.

upvoted 2 times

🗨️ **mghectorenjoyer69** 10 months, 1 week ago

c ra unga amma

upvoted 1 times

🗨️ **ubiqinon** 1 year, 8 months ago

B is the least overhead solution

upvoted 3 times

🗨️ 👤 **zk1200** 1 year, 9 months ago

Selected Answer: B

I considered D as well which refers to using CodeDeploy. however using codedeploy adds more work. So alias makes more sense.

upvoted 2 times

🗨️ 👤 **ihta_2031** 1 year, 10 months ago

Selected Answer: B

lambda function version => alias

upvoted 4 times

🗨️ 👤 **aragon_saa** 1 year, 10 months ago

B

<https://www.examttopics.com/discussions/amazon/view/96149-exam-aws-certified-developer-associate-topic-1-question-441/>

upvoted 3 times

A developer has written an AWS Lambda function. The function is CPU-bound. The developer wants to ensure that the function returns responses quickly.

How can the developer improve the function's performance?

- A. Increase the function's CPU core count.
- B. Increase the function's memory.
- C. Increase the function's reserved concurrency.
- D. Increase the function's timeout.

Correct Answer: B

Community vote distribution



🗳️ **ihnta_2031** Highly Voted 1 year, 10 months ago

Selected Answer: B

Cpu utilisation => increase memory

upvoted 15 times

🗳️ **Kashan6109** Highly Voted 1 year, 6 months ago

Selected Answer: B

Option B is correct, the only adjustable parameter (in terms of hardware) is lambda memory. Increasing lambda memory will result in automatic adjustment of CPU.

Lambda memory is adjustable from 128 MB upto 10 GB

upvoted 9 times

🗳️ **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B

A) Eliminated - We cannot directly control the CPU core count.

B) Correct - increasing memory not only gives more memory but also allocates more CPU power to the function.

C) Eliminated - Reserved concurrency controls how many instances of the function can run at the same time, but it doesn't improve the performance of a single function execution

D) Eliminated - Increasing the timeout allows the function to run longer, but it doesn't improve the speed of the function

upvoted 1 times

🗳️ **trieudo** 1 month, 2 weeks ago

Selected Answer: B

==> discard A: can't directly increase lambda CPU

==> discard C: concurrecny for handling many tasks at same time capability, not increase speed response when CPU-bound

=> discard D: be not relevant, increase maximum time to wait function to finish

B: with lambda, you can increase memory, implicity CPU usage for lambda also up too

upvoted 1 times

🗳️ **rue_** 3 months ago

Selected Answer: B

the only way to adjust the vCPU assigned to your function, is through an increase in memory

upvoted 1 times

🗳️ **tgood** 5 months, 3 weeks ago

Selected Answer: B

B is correct

upvoted 1 times

🗳️ **65703c1** 8 months, 2 weeks ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗨️ **leonardoliveros** 1 year, 2 months ago

Selected Answer: B

If you increase the memory on a Lambda Function hence your vCPU also increased

upvoted 1 times

🗨️ **james2033** 1 year, 2 months ago

Selected Answer: B

Quote 'If a function is CPU-, network- or memory-bound, then changing the memory setting can dramatically improve its performance.' at

<https://docs.aws.amazon.com/lambda/latest/operatorguide/computing-power.html>

upvoted 3 times

🗨️ **Majong** 1 year, 8 months ago

Selected Answer: B

Lambda allocates CPU power in proportion to the amount of memory configured. You can read more here:

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-function-common.html#configuration-memory-console>

upvoted 5 times

🗨️ **Devon_Fazekas** 1 year, 8 months ago

Increasing the function's CPU core count is not an option in AWS Lambda. AWS Lambda automatically manages the allocation of CPU power and only allows scaling of memory.

upvoted 2 times

🗨️ **geekdamsel** 1 year, 8 months ago

Correct answer is B.

upvoted 3 times

🗨️ **Bibay** 1 year, 9 months ago

Selected Answer: B

. Increase the function's memory.

The performance of an AWS Lambda function is primarily determined by the amount of allocated memory. When you increase the memory, you also increase the available CPU and network resources. This can result in faster execution times, especially for CPU-bound functions. Increasing the CPU core count, reserved concurrency, or timeout may not have as significant an impact on performance as increasing memory.

upvoted 1 times

🗨️ **blathul** 1 year, 9 months ago

Selected Answer: B

Adding more memory proportionally increases the amount of CPU, increasing the overall computational power available. If a function is CPU-, network- or memory-bound, then changing the memory setting can dramatically improve its performance.

<https://docs.aws.amazon.com/lambda/latest/operatorguide/computing-power.html>

upvoted 1 times

🗨️ **Syre** 1 year, 9 months ago

Selected Answer: A

On this particular question the answer is A.

while increasing memory can indirectly improve CPU performance, it's not always the most effective solution for CPU-bound functions, and increasing the CPU core count is usually a better option for improving performance in such cases. Please note - CPU-Bound functions. This question is to trick you

upvoted 1 times

🗨️ **Majong** 1 year, 8 months ago

In this particular question it is B. You are right that in normal question it might be A but for a Lambda function you are not able to change the CPU. Lambda allocates CPU power in proportion to the amount of memory configured. You can read more here:

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-function-common.html#configuration-memory-console>

upvoted 4 times

🗨️ **Untamables** 1 year, 10 months ago

Selected Answer: B

B

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-function-common.html#configuration-memory-console>
upvoted 3 times

For a deployment using AWS Code Deploy, what is the run order of the hooks for in-place deployments?

- A. BeforeInstall -> ApplicationStop -> ApplicationStart -> AfterInstall
- B. ApplicationStop -> BeforeInstall -> AfterInstall -> ApplicationStart
- C. BeforeInstall -> ApplicationStop -> ValidateService -> ApplicationStart
- D. ApplicationStop -> BeforeInstall -> ValidateService -> ApplicationStart

Correct Answer: B

Community vote distribution

B (88%) 12%

 **pratchatcap** Highly Voted 1 year, 10 months ago


Selected Answer: B

It's B. Check the image in the link.

<https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html#appspec-hooks-server>
upvoted 24 times

 **awsdummie** 1 year, 8 months ago

Answer A For InPlace deployment
upvoted 2 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B

ApplicationStop: Stops the currently running application (if any) on the instance.
BeforeInstall: Runs tasks or scripts before the new application revision is installed.
Install: Deploys the new application revision.
AfterInstall: Runs tasks or scripts after the new application revision is installed.
ApplicationStart: Starts the application after installation.
ValidateService: Validates that the service is running as expected after the application starts.
upvoted 1 times


 **sumanshu** 1 month, 1 week ago

A) Eliminated - The ApplicationStop hook must be the first step to stop the currently running application

C) BeforeInstall cannot occur before ApplicationStop because the application needs to be stopped first.

D) Eliminated - ValidateService should run after the application starts, not before.


upvoted 1 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: B


==> discard C, D: ValidateService is invalid lifecycle in In-place CodeDeploy
==> discard A: it is not logic, cuz install --> then stop app ==> pointless

So I choose B
upvoted 1 times

 **Anandesh** 6 months, 2 weeks ago

Selected Answer: B

<https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html#appspec-hooks-server>
upvoted 1 times

 **65703c1** 8 months, 2 weeks ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

SD_CS 12 months ago

Selected Answer: B

Answer is B. There is no doubt - please go to the URL <https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html> and search with "In-place deployments"

In fact none of the deployments follow the order mentioned in A

upvoted 2 times

ez_24 1 year, 1 month ago

B

In AWS CodeDeploy for in-place deployments, the hooks run in the following order:

ApplicationStop: Executed before the new application revision is downloaded.

DownloadBundle: The new application revision is downloaded.

BeforeInstall: Executed after the new revision is downloaded but before the new version is installed.

Install: The application revision specified in the deployment is installed.

AfterInstall: Executed after the application revision is installed.

ApplicationStart: Invoked to start any services that were stopped during ApplicationStop.

ValidateService: Ensures the service is operating correctly after the new deployment.

This sequence ensures a smooth deployment process by systematically stopping, updating, and restarting the application.

upvoted 4 times

quanbui 1 year, 3 months ago

ApplicationStop -> BeforeInstall -> AfterInstall -> ApplicationStart -> ValidateService.

Ref: <https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html>

upvoted 1 times

Skywalker23 1 year, 4 months ago

Selected Answer: B

Application must be stopped before installation. Otherwise the installation may corrupt the running application's files and cause damages. Not good.

upvoted 2 times

Tony88 1 year, 5 months ago

Selected Answer: B

Stopped -> Installed -> Started -> Validated

Go with B.

upvoted 3 times

ninomfr64 1 year, 5 months ago

Selected Answer: B

It's B as per doc <https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html#appspec-hooks-server-~:text=a%20load%20balancer-.,Lifecycle%20event%20hook%20availability,-The%20following%20table>

upvoted 1 times

sp323 1 year, 5 months ago

Application start is after install

upvoted 1 times

fcbc62d 1 year, 6 months ago

Selected Answer: B

For in-place deployment B is correct.

<https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html>

upvoted 1 times

jipark 1 year, 6 months ago

Selected Answer: B

this image explain all :

<https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html#appspec-hooks-server>

upvoted 1 times

🗨️ 👤 **ScherbakovMike** 1 year, 8 months ago

Definitely, B: the order is the same in case of InPlace and Blue/Green deployment:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html#reference-appspec-file-structure-hooks-availability>

upvoted 1 times

🗨️ 👤 **Nagendhar** 1 year, 8 months ago

Ans: A

For an in-place deployment using AWS CodeDeploy, the run order of the hooks is option A, "BeforeInstall -> ApplicationStop -> ApplicationStart -> AfterInstall."

This is the correct order of hooks for an in-place deployment, where the deployment package is installed on the same set of Amazon EC2 instances that are running the current version of the application.

upvoted 2 times

🗨️ 👤 **DeaconStJohn** 1 year, 9 months ago

Selected Answer: B

I'll go with B based on the link provided by others

upvoted 2 times

A company is building a serverless application on AWS. The application uses an AWS Lambda function to process customer orders 24 hours a day, 7 days a week. The Lambda function calls an external vendor's HTTP API to process payments.

During load tests, a developer discovers that the external vendor payment processing API occasionally times out and returns errors. The company expects that some payment processing API calls will return errors.

The company wants the support team to receive notifications in near real time only when the payment processing external API error rate exceed 5% of the total number of transactions in an hour. Developers need to use an existing Amazon Simple Notification Service (Amazon SNS) topic that is configured to notify the support team.

Which solution will meet these requirements?

- A. Write the results of payment processing API calls to Amazon CloudWatch. Use Amazon CloudWatch Logs Insights to query the CloudWatch logs. Schedule the Lambda function to check the CloudWatch logs and notify the existing SNS topic.
- B. Publish custom metrics to CloudWatch that record the failures of the external payment processing API calls. Configure a CloudWatch alarm to notify the existing SNS topic when error rate exceeds the specified rate.
- C. Publish the results of the external payment processing API calls to a new Amazon SNS topic. Subscribe the support team members to the new SNS topic.
- D. Write the results of the external payment processing API calls to Amazon S3. Schedule an Amazon Athena query to run at regular intervals. Configure Athena to send notifications to the existing SNS topic when the error rate exceeds the specified rate.

Correct Answer: B

Community vote distribution

B (100%)

 **Bibay** Highly Voted 1 year, 9 months ago


Selected Answer: B

B. Publish custom metrics to CloudWatch that record the failures of the external payment processing API calls. Configure a CloudWatch alarm to notify the existing SNS topic when the error rate exceeds the specified rate is the best solution to meet the requirements.

With CloudWatch custom metrics, developers can publish and monitor custom data points, including the number of failed requests to the external payment processing API. A CloudWatch alarm can be configured to notify an SNS topic when the error rate exceeds the specified rate, allowing the support team to be notified in near real-time.

Option A is not optimal since it involves scheduling a Lambda function to check the CloudWatch logs. Option C may not provide the desired functionality since it does not specify a rate at which to notify the support team. Option D is more complex than necessary, as it involves writing the results to S3 and configuring an Athena query to send notifications to an SNS topic.

upvoted 14 times

 **Untamables** Highly Voted 1 year, 10 months ago


Selected Answer: B

The correct answer is B.

You can use the Embedded Metrics format to embed custom metrics alongside detailed log event data. CloudWatch automatically extracts the custom metrics so you can visualize and alarm on them, for real-time incident detection.

<https://docs.aws.amazon.com/lambda/latest/operatorguide/custom-metrics.html>

upvoted 6 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B

A) Eliminated - Requires manual log queries and a custom Lambda function to process logs, introducing operational overhead.

B) Correct - CloudWatch Alarms can calculate error rates using metric math and automatically notify the SNS topic.

C) Eliminated - This approach notifies the support team for every API failure, not just when the error rate exceeds 5%.

D) Eliminated - Adds significant operational complexity (managing S3 storage, Athena queries, and scheduling).

upvoted 1 times

🗨️ **trieudo** 1 month, 2 weeks ago

Selected Answer: B

keyword: error rate exceed 5%, receive notifications in near real time

==> discard A: use lambda or athena with schedule, violate ' receive notifications in near real time '

==> discard C: die in spam of message if any, and don't know when error up 5%

B is most good solution: CloudWatch Alarm can calculate 5% with setup, and notify righ away when reach thresold set-up

upvoted 1 times

🗨️ **rue_** 3 months ago

Selected Answer: B

It's B, allows you to customize the metrics as required in the question, and sends notification in near real time instead of polling

upvoted 1 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗨️ **Tony88** 1 year, 5 months ago

Selected Answer: B

Require "near real-time" notification, so you should not use scheduled solution.

Creating a new SNS topic is no sense.

upvoted 2 times

🗨️ **Ponyi** 1 year, 2 months ago

In the question, it is also mentioned that "Developer needs to use the existing SNS topic...."

upvoted 1 times

🗨️ **jayvarma** 1 year, 5 months ago

Option B. Using custom metrics, Developers will be able to publish and monitor custom data points such as the no. of failed requests to the external payment processing API. Create a CloudWatch alarm and configure it to be triggered when the rate of error exceeds the specified number in the question.

upvoted 1 times

🗨️ **svrnvtr** 1 year, 10 months ago

Selected Answer: B

It is B

upvoted 3 times

A company is offering APIs as a service over the internet to provide unauthenticated read access to statistical information that is updated daily. The company uses Amazon API Gateway and AWS Lambda to develop the APIs. The service has become popular, and the company wants to enhance the responsiveness of the APIs.

Which action can help the company achieve this goal?

- A. Enable API caching in API Gateway.
- B. Configure API Gateway to use an interface VPC endpoint.
- C. Enable cross-origin resource sharing (CORS) for the APIs.
- D. Configure usage plans and API keys in API Gateway.

Correct Answer: A

Community vote distribution

A (100%)

🗳️ **Bibay** Highly Voted 1 year, 9 months ago

Selected Answer: A

A. Enable API caching in API Gateway can help the company enhance the responsiveness of the APIs. By enabling caching, API Gateway stores the responses from the API and returns them for subsequent requests instead of forwarding the requests to Lambda. This reduces the number of requests to Lambda, improves API performance, and reduces latency for users.

upvoted 29 times

🗳️ **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: A

A) Cached data reduces calls to the backend (in this case, Lambda) and serves responses directly from the cache. Since the data updates only once per day, caching is highly effective

B) Eliminated - This is useful for securing private APIs inside a VPC, but the scenario involves public, unauthenticated APIs over the internet

C) Eliminated - CORS is related to resolving client-side browser compatibility issues

D) Eliminated - Usage plans and API keys manage access control and rate limiting, which are useful for monetization

upvoted 1 times

🗳️ **trieudo** 1 month, 2 weeks ago

Selected Answer: A

keyword: updated daily, enhance the responsiveness

=> A

upvoted 1 times

🗳️ **ahadh7621** 6 months, 1 week ago

Selected Answer: A

Answer is A

upvoted 1 times

🗳️ **65703c1** 8 months, 2 weeks ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗳️ **badsati** 9 months, 3 weeks ago

Selected Answer: A

Answer is A. Caching will enhance the responsiveness of the APIs.

upvoted 1 times

🗳️ **leonardoliveros** 1 year, 2 months ago

Selected Answer: A

Caching the request is the best option because the request don't forwarded to Lambda Function and this reduces latency and also recude costs
upvoted 2 times

🗨️ **Tony88** 1 year, 5 months ago

Selected Answer: A

Go with A.

A. Caching is the general solution to improve performance of non-frequently change data. (in this case, daily, not really frequent)

B. interface endpoint is a VPC concept, in this architect we don't need to concern with VPC. For those who are interested, go check with interface endpoint and gateway endpoint.

C. CORS is short for cross origin resource share. it is a distractor here. You may consider CORS when your client cannot access to your API Gateway resource, not when you want to improve the performance.

D. usage plan is used when your API client's behaviour is predictable, and it can avoid anormal usage.

upvoted 4 times

🗨️ **yruryenucakc** 1 year, 5 months ago

A-> Caching frequently accessed api calls allows reducing process time every time api is called.

B-> You shloud configure VPC if you want to change network security of your application. So it does not neccessarily increase the performance.

C-> CORS (Cross Origin Resource Sharing), allows you to proccess the api calls that comes from outside of your AWS organization.

Again nothing to do with the performance. One of the use case of this feature is if you want to keep your web app apis reachable from public internet you should enable CORS for it.

D → This is mainly for throttling and controlling who can access the API and at what rate. While it's useful for controlling and metering access, it doesn't enhance the responsiveness of the API

upvoted 1 times

🗨️ **svrnvtr** 1 year, 10 months ago

Selected Answer: A

I vote for A

upvoted 3 times

🗨️ **Untamables** 1 year, 10 months ago

Selected Answer: A

A

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-caching.html>

upvoted 3 times

A developer wants to store information about movies. Each movie has a title, release year, and genre. The movie information also can include additional properties about the cast and production crew. This additional information is inconsistent across movies. For example, one movie might have an assistant director, and another movie might have an animal trainer.

The developer needs to implement a solution to support the following use cases:

For a given title and release year, get all details about the movie that has that title and release year.

For a given title, get all details about all movies that have that title.

For a given genre, get all details about all movies in that genre.

Which data store configuration will meet these requirements?

A. Create an Amazon DynamoDB table. Configure the table with a primary key that consists of the title as the partition key and the release year as the sort key. Create a global secondary index that uses the genre as the partition key and the title as the sort key.

B. Create an Amazon DynamoDB table. Configure the table with a primary key that consists of the genre as the partition key and the release year as the sort key. Create a global secondary index that uses the title as the partition key.

C. On an Amazon RDS DB instance, create a table that contains columns for title, release year, and genre. Configure the title as the primary key.

D. On an Amazon RDS DB instance, create a table where the primary key is the title and all other data is encoded into JSON format as one additional column.

Correct Answer: A

Community vote distribution

A (100%)


 **Bibay** Highly Voted 1 year, 9 months ago

Selected Answer: A

A. Create an Amazon DynamoDB table. Configure the table with a primary key that consists of the title as the partition key and the release year as the sort key. Create a global secondary index that uses the genre as the partition key and the title as the sort key.

This option is the best choice for the given requirements. By using DynamoDB, the developer can store the movie information in a flexible and scalable NoSQL database. The primary key can be set to the title and release year, allowing for efficient retrieval of information about a specific movie. The global secondary index can be created using the genre as the partition key, allowing for efficient retrieval of information about all movies in a specific genre. Additionally, the use of a NoSQL database like DynamoDB allows for the flexible storage of additional properties about the cast and crew, as each movie can have different properties without affecting the structure of the database.

upvoted 16 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: A

B) Eliminated - No efficient way to query by title + release year since the primary key is genre + release year, and the GSI only has the title

C) Eliminated - Adding new columns for properties like assistant director or animal trainer becomes cumbersome and violates flexibility.

D) Eliminated - JSON in RDS is harder to query

upvoted 1 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: A

keyword: (title + release year), title, genre, each film has differential cast and crew

=> Discard C,D: when it store film by structure && option C must scan all data that is not primary key(title), seem be not flexible && option D is horrible way, when storing data as json structure, take time to extract this json

=> discard B: query (title + release year) must scan all item without combining them in primary key(partition key and sort key) at primary table or GSI

upvoted 1 times

 **ahadh7621** 6 months, 1 week ago

Selected Answer: A

This question was on my exam July 23rd, 2024. Answer is A

upvoted 2 times

🗨️ **sid4510** 7 months ago

A is mostly correct, but I do see one problem there because in one year there can be same title movies can come which invalidate our primary key having title as partition key and year as a. Sort. key

upvoted 1 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ **leonardoliveros** 1 year, 2 months ago

Selected Answer: A

If you create a primary key with title(pk) and release(sk) date you covered two scenarios, and also you need a GSI by last scenario with genre so you should creating a GSI with genre (pk) and title (sk)

upvoted 2 times

🗨️ **Tony88** 1 year, 5 months ago

Selected Answer: A

Go with A.

NoSQL is good when data attributes are inconsistent -> DynamoDB

Primary key should be unique, go with title + release year.

upvoted 4 times

🗨️ **jayvarma** 1 year, 5 months ago

As the schema for each entry of data into the database is not the same all the time, We would require a NoSQL database. So, RDS DB instance is ruled out. The answer is between A and B.

As we would need the partition key to be as unique as possible, we would like to have the title of the movie as the partition key. Because having the partition key as the genre will create a hot partition problem and our data stored in the DynamoDB will be skewed.

So option A is the answer.

upvoted 4 times

🗨️ **Krok** 1 year, 10 months ago

Selected Answer: A

It's A - I totally agree. It's a single appropriate solution. But in my opinion genre isn't a quite good option as GSI partition key - it isn't high distribution and we can get a hot partition.

upvoted 2 times

🗨️ **shahs10** 1 year, 10 months ago

Selected Answer: A

Option A because we have to search on the basis of title so it is better to partition by title. Also we have to search by genre so it is good option to make GSI using genre as partition key

upvoted 2 times

🗨️ **Untamables** 1 year, 10 months ago

Selected Answer: A

The correct answer is A.

Amazon DynamoDB is suited for storing inconsistent attributes data across items.

Option B is wrong. This solution does not help get items with the condition of the combination, title and release year.

upvoted 3 times

A developer maintains an Amazon API Gateway REST API. Customers use the API through a frontend UI and Amazon Cognito authentication. The developer has a new version of the API that contains new endpoints and backward-incompatible interface changes. The developer needs to provide beta access to other developers on the team without affecting customers. Which solution will meet these requirements with the LEAST operational overhead?

- A. Define a development stage on the API Gateway API. Instruct the other developers to point the endpoints to the development stage.
- B. Define a new API Gateway API that points to the new API application code. Instruct the other developers to point the endpoints to the new API.
- C. Implement a query parameter in the API application code that determines which code version to call.
- D. Specify new API Gateway endpoints for the API endpoints that the developer wants to add.

Correct Answer: A

Community vote distribution

A (83%)

B (17%)

 **Bibay** Highly Voted 1 year, 9 months ago

Selected Answer: A

Option A is the correct solution to meet the requirements with the least operational overhead.

Defining a development stage on the API Gateway API enables other developers to test the new version of the API without affecting the production environment. This approach allows the developers to work on the new version of the API independently and avoid conflicts with the production environment.


The other options involve creating a new API or new endpoints, which could introduce additional operational overhead, such as managing multiple APIs or endpoints, configuring access control, and updating the frontend UI to point to the new endpoints or API. Option C also introduces additional complexity by requiring the implementation of a query parameter to determine which code version to call.

upvoted 12 times

 **beekeeper0101** 7 months, 3 weeks ago

Thank you!

upvoted 1 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: A


A) Correct - The developer can deploy the new version of the API to a development stage while keeping the existing API in the production stage for customers. Beta users (team developers) can use the unique URL of the development stage to access the new API version.

B) Eliminated - This involves duplicating the API Gateway setup and potentially managing additional resources

C) Eliminated - This requires modifying the application logic, which increases the development effort and complexity. - Add complexity and operational overhead.

D) Eliminated - The goal is to let beta users test the new API without affecting existing customers. Mixing endpoints in the same API Gateway does not provide a clean separation, which increases the risk of customer impact.

upvoted 1 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: A

keyword: LEAST operational overhead, backward-incompatible, beta access to other developers on the team without affecting customers. (separate)

==> Discard B: take effort to operation and maintain many API gateway instances

==> Discard C: It is not separate API production and beta, make effort to understand, transfer, or fixing something when misunderstanding. Violate backward-incompatible, when new version can also use old endpoint version 100% (just adding, not delete or modify)

==> Discard D: like a bit with C, mixing 2 environment prod and beta,

upvoted 1 times

🗨️ **AnthonyTL** 4 months, 2 weeks ago

A is correct as the Q asked for 'with the LEAST operational overhead'

C got less operation than B but more than A

upvoted 1 times

🗨️ **ACurryDeveloper** 6 months ago

Selected Answer: A

It is A you chutiyas

upvoted 1 times

🗨️ **ahadh7621** 6 months, 1 week ago

This question was on my exam July 23rd, 2024. Answer is A

upvoted 2 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ **mghectorenjoyer69** 10 months, 1 week ago

Selected Answer: B

B ra unga amma

upvoted 1 times

🗨️ **HayLLIHuK** 10 months, 2 weeks ago

Selected Answer: A

The correct answer is A (development stage).

You can configure a development stage for your API Gateway API and then integrate it with the new version of the backend functionality that has new endpoints and backward-incompatible interface changes. The customers can continue to use the existing API.

upvoted 1 times

🗨️ **41eb566** 10 months, 3 weeks ago

Selected Answer: B

The option that meets the requirements with the least operational overhead is:

B. Define a new API Gateway API that points to the new API application code. Instruct the other developers to point the endpoints to the new API.

Here's why:

A. Defining a development stage on the existing API Gateway API could potentially affect customers if not managed properly. It might introduce changes or issues to the existing API that customers are using.

C. Implementing a query parameter in the API application code to determine the code version introduces complexity and potential risk, as it requires changes to the application code itself. It also doesn't isolate the beta access from the main API.

D. Specifying new API Gateway endpoints for the new API endpoints adds complexity and overhead. It requires managing multiple endpoints, potentially affecting the API's simplicity and increasing maintenance overhead.

upvoted 2 times

🗨️ **hungnv6_rikkei** 11 months, 4 weeks ago

A is answer

upvoted 1 times

🗨️ **Aleam** 1 year, 1 month ago

Selected Answer: B

LEAST operational overhead would be B.

upvoted 2 times

🗨️ **leonardoliveros** 1 year, 2 months ago

Selected Answer: A

The stages gives the capacity to tests a new version in an APiG without affecting customers in others stages

upvoted 2 times

🗨️ 👤 **Tony88** 1 year, 5 months ago

Selected Answer: A

The best practice is to define a development stage.

upvoted 3 times

🗨️ 👤 **jayvarma** 1 year, 5 months ago

Option A is the right answer. Defining a development stage on the API Gateway API would provide other developers with a way to test the newer version of the API without affecting prod.

The rest of the options would create a lot of operational overhead.

upvoted 1 times

🗨️ 👤 **MrTee** 1 year, 9 months ago

Selected Answer: A

The developer should define a development stage on the API Gateway API. They should then instruct the other developers to point the endpoints to the development stage. This solution will meet the requirements with the least operational overhead

upvoted 1 times

🗨️ 👤 **Untamables** 1 year, 10 months ago

Selected Answer: A

A

<https://docs.aws.amazon.com/apigateway/latest/developerguide/set-up-stages.html>

<https://docs.aws.amazon.com/apigateway/latest/developerguide/canary-release.html>

upvoted 3 times

A developer is creating an application that will store personal health information (PHI). The PHI needs to be encrypted at all times. An encrypted Amazon RDS for MySQL DB instance is storing the data. The developer wants to increase the performance of the application by caching frequently accessed data while adding the ability to sort or rank the cached datasets.


Which solution will meet these requirements?

- A. Create an Amazon ElastiCache for Redis instance. Enable encryption of data in transit and at rest. Store frequently accessed data in the cache.
- B. Create an Amazon ElastiCache for Memcached instance. Enable encryption of data in transit and at rest. Store frequently accessed data in the cache.
- C. Create an Amazon RDS for MySQL read replica. Connect to the read replica by using SSL. Configure the read replica to store frequently accessed data.
- D. Create an Amazon DynamoDB table and a DynamoDB Accelerator (DAX) cluster for the table. Store frequently accessed data in the DynamoDB table.

Correct Answer: A

Community vote distribution

A (100%)

 **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: A


A

You can use Amazon ElastiCache for Redis Sorted Sets to easily implement a dashboard that keeps a list of sorted data by their rank.

<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/elasticache-use-cases.html#elasticache-for-redis-use-cases-gaming>

<https://aws.amazon.com/elasticache/redis-vs-memcached/>

upvoted 17 times

 **jipark** 1 year, 6 months ago

in sum,

REDIS featured encryption, PCI-DSS

MemCache support AutoDiscovery

upvoted 4 times

 **AnthonyTL** 4 months, 2 weeks ago

Both supported encryption, but MemCache doesn't support 'sort or rank'

upvoted 2 times

 **Bibay** Highly Voted 1 year, 9 months ago

Selected Answer: A

To meet the requirements of caching frequently accessed data while adding the ability to sort or rank cached datasets, a developer should choose Amazon ElastiCache for Redis. ElastiCache is a web service that provides an in-memory data store in the cloud, and it supports both Memcached and Redis engines. While both engines are suitable for caching frequently accessed data, Redis is a better choice for this use case because it provides sorted sets and other data structures that allow for sorting and ranking of cached datasets. The data in ElastiCache can be encrypted at rest and in transit, ensuring the security of the PHI. Therefore, option A is the correct answer.

upvoted 12 times

 **Hasitha99** Most Recent 3 weeks ago

Selected Answer: A

* Amazon ElastiCache service provide in-memory cache in the cloud.

* Redis and MemCache are the 2 popular options which ElastiCache offerd for us.

* Both support encryption at rest and transist (previously only redis supported this).

* Redis offerd advance features like datastructures, transactions ,pub/sub etc.

Since question is asking about most cost effective solution(already buit with MYSQL) , we have to ignore the DynamoDB option .

Now we can pick Redis due to its natual sort facility.

upvoted 1 times

🗉 👤 **sumanshu** 1 month, 1 week ago

Selected Answer: A

B) Eliminated - Memcached does not support advanced data structure

C) Eliminated - A read replica is not a true caching solution. It is a secondary database instance that replicates data for read-heavy workloads

D) Eliminated - DynamoDB with DAX adds more operational overhead because you would be introducing a NoSQL database into the architecture when you're already using Amazon RDS (relational database) for storing data.

upvoted 1 times

🗉 👤 **trieudo** 1 month, 2 weeks ago

Selected Answer: A

keyword: caching, sort, rank

==> discard B: not support 'sort' and 'rank'

==> discard C: not 'caching'

==> discard D: caching only support NoSQL dynamoDB && sort and rank belong to primary key (partition and secondary) of DynamoDB

A fit all requirement

upvoted 1 times

🗉 👤 **ahadh7621** 6 months, 1 week ago

Selected Answer: A

This question was on my exam July 23rd, 2024. Answer is A.

upvoted 4 times

🗉 👤 **tsangckl** 7 months, 2 weeks ago

This appear at 17 Jun exam

upvoted 4 times

🗉 👤 **65703c1** 8 months, 2 weeks ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗉 👤 **Vaibs099** 9 months, 1 week ago

Sorted Sets in ElastiCache redis can do this job.

upvoted 1 times

🗉 👤 **leonardoliveros** 1 year, 2 months ago

Selected Answer: A

Redis is the best option to cached the results of queries and it also offer a encryption in-transit and at-rest

upvoted 1 times

🗉 👤 **nmc12** 1 year, 4 months ago

Redis: Supports various data structures such as strings, hashes, lists, sets, sorted sets, bitmaps, hyperloglogs, and geospatial indexes.

Memcached: Primarily supports string-based keys and values; does not support advanced data structures.

upvoted 4 times

🗉 👤 **brandon87** 1 year, 10 months ago

Selected Answer: A

ElastiCache for Redis also features Online Cluster Resizing, supports encryption, and is HIPAA eligible and PCI DSS compliant.

<https://aws.amazon.com/elasticache/redis-vs-memcached/>

upvoted 7 times

A company has a multi-node Windows legacy application that runs on premises. The application uses a network shared folder as a centralized configuration repository to store configuration files in .xml format. The company is migrating the application to Amazon EC2 instances. As part of the migration to AWS, a developer must identify a solution that provides high availability for the repository. Which solution will meet this requirement MOST cost-effectively?

- A. Mount an Amazon Elastic Block Store (Amazon EBS) volume onto one of the EC2 instances. Deploy a file system on the EBS volume. Use the host operating system to share a folder. Update the application code to read and write configuration files from the shared folder.
- B. Deploy a micro EC2 instance with an instance store volume. Use the host operating system to share a folder. Update the application code to read and write configuration files from the shared folder.
- C. Create an Amazon S3 bucket to host the repository. Migrate the existing .xml files to the S3 bucket. Update the application code to use the AWS SDK to read and write configuration files from Amazon S3.
- D. Create an Amazon S3 bucket to host the repository. Migrate the existing .xml files to the S3 bucket. Mount the S3 bucket to the EC2 instances as a local volume. Update the application code to read and write configuration files from the disk.

Correct Answer: C

Community vote distribution

C (75%)

D (25%)

shahs10 Highly Voted 1 year, 10 months ago

Why is not there EFS to replace shared file system
upvoted 15 times

albert_kuo 4 months, 2 weeks ago

Windows legacy application
upvoted 2 times

[Removed] 1 year, 1 month ago

This is what I was looking for - but not an option
upvoted 2 times

alexandru_tata 10 months, 1 week ago

EFS would not have been an option. It does not work for Linux:
<https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/AmazonEFS.html>
upvoted 4 times

serverlessme 3 months, 1 week ago

It does not work for Windows you meant to say?
upvoted 3 times

nmc12 1 year, 4 months ago

it is best solution. But we can use S3 without EFS
upvoted 3 times

Bibay Highly Voted 1 year, 8 months ago

c

Option C is the most cost-effective solution to provide high availability for the centralized configuration repository. Amazon S3 provides a highly durable and available object storage service. S3 stores objects redundantly across multiple devices and multiple facilities within a region, making it highly available. The developer can migrate the existing .xml files to an S3 bucket and update the application code to use the AWS SDK to read and write configuration files from Amazon S3.

Option A and B are not the best solutions as they require the developer to use the host operating system to share a folder, which can lead to a single point of failure.

Option D is not a recommended solution as it is not a direct way of accessing an S3 bucket. While it is possible to use third-party tools to mount an S3 bucket as a local disk, it can lead to performance issues and additional complexity.

upvoted 11 times

🗨️ **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: C

A) Eliminated - We need to deploy and maintain a file system on EBS, which could lead to additional complexity

B) Eliminated - Instance store volumes are ephemeral and will be lost if the EC2 instance is stopped or terminated.

D) Eliminated - Mounting S3 directly as a volume is not natively supported without third-party solutions or tools like s3fs, which could add unnecessary complexity and overhead.

upvoted 2 times

🗨️ **trieudo** 1 month, 2 weeks ago

Selected Answer: C

Keyword: multi-node, Windows legacy application, shared folder as a centralized configuration repository, high availability, MOST cost-effectively

maybe many of you think EFS first when reach 'shared folder', but EFS is not work with 'Windows legacy application', it depends on OS, MOST cost-effectively

==> discard B: when you reboot or turn off EC2, data in 'instance store volume' will be lost, violate 'high availability'

==> discard A: EBS at time, can link only one EC2, violate 'shared folder'

==> discard D: complexity than option C, you must setup more, then maintain many things afterward, violate 'MOST cost-effectively'

C is match with all requirements

upvoted 2 times

🗨️ **f271c23** 1 month, 3 weeks ago

Selected Answer: D

I am not sure why D is not the best option. This could give a no code solution to legacy application.

upvoted 1 times

🗨️ **Venky786** 2 months, 3 weeks ago

Why A is not correct answer ?

upvoted 1 times

🗨️ **AWSArt** 1 month, 1 week ago

Because EBS is way more costly than s3

upvoted 1 times

🗨️ **tgood** 5 months, 3 weeks ago

Selected Answer: C

i think EFS should be added to answer

upvoted 1 times

🗨️ **serverlessme** 3 months, 1 week ago

EFS does not work for Windows.

upvoted 1 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗨️ **teddyjr** 9 months, 1 week ago

S3 mountpoint does not support editing files or deleting directories.

So the answer is C

upvoted 2 times

🗨️ **Dikshika** 9 months, 2 weeks ago

D since we have s3 mountpoint available now

<https://aws.amazon.com/about-aws/whats-new/2023/03/mountpoint-amazon-s3/>

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/mountpoint.html>

upvoted 1 times

🗨️ **someone234** 11 months, 3 weeks ago

Selected Answer: C

Option C is the most cost-effective solution to provide high availability for the centralized configuration repository. Amazon S3 provides a highly durable and available object storage service. S3 stores objects redundantly across multiple devices and multiple facilities within a region, making it highly available. The developer can migrate the existing .xml files to an S3 bucket and update the application code to use the AWS SDK to read and write configuration files from Amazon S3.

upvoted 2 times

🗨️ 👤 **gqs3119** 1 year, 1 month ago

Selected Answer: D

Today It's D.

Few months ago I'd pick C, but since then amazon released mountpoint for linux, so it's possible to mount S3 on any major Linux distro, by using WSL 2 it is also possible to mount S3 on Windows. Doing so cuts the cost of modifying the legacy application.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/mountpoint-installation.html>

<https://aws.plainenglish.io/mounting-amazon-s3-buckets-on-windows-52b5f1434cd7>

upvoted 2 times

🗨️ 👤 **letsLearn_001** 2 months, 1 week ago

Mount doesn't support for Windows.

upvoted 1 times

🗨️ 👤 **squeeze_talus0y** 1 year ago

Your solution overcomplicates things.

upvoted 3 times

🗨️ 👤 **SD_CS** 12 months ago

But the apps are legacy windows app so mountpoints will not help - my opinion

upvoted 1 times

🗨️ 👤 **leonardoliveros** 1 year, 2 months ago

Selected Answer: C

EBS and Instance Store just attached one instance so these's expense and don't scalable, and S3 it's the best option to handle the repository of .xml because it's very scalable and low-cost

upvoted 2 times

🗨️ 👤 **HanTran0795** 1 year, 3 months ago

Selected Answer: D

It is a Windows legacy application. What if the sdk doesn't support the app? I choose D.

upvoted 2 times

🗨️ 👤 **ronn555** 1 year, 2 months ago

C

S3 Buckets can only be mounted directly to Linux EC2 instances

upvoted 1 times

🗨️ 👤 **gqs3119** 1 year, 1 month ago

It can be mounted to many distros today, and using WSL2 also to Windows.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/mountpoint-installation.html>

upvoted 2 times

🗨️ 👤 **AhmedAliHashmi** 1 year, 5 months ago

Correct answer is C

upvoted 1 times

🗨️ 👤 **senadevtrd** 1 year, 8 months ago

Selected Answer: C

In these options, this is more correct

upvoted 1 times

🗨️ 👤 **Untamables** 1 year, 10 months ago

Selected Answer: C

C

<https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/AmazonS3.html>

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingAWSSDK.html>

upvoted 6 times

A company wants to deploy and maintain static websites on AWS. Each website's source code is hosted in one of several version control systems, including AWS CodeCommit, Bitbucket, and GitHub.

The company wants to implement phased releases by using development, staging, user acceptance testing, and production environments in the AWS Cloud. Deployments to each environment must be started by code merges on the relevant Git branch. The company wants to use HTTPS for all data exchange. The company needs a solution that does not require servers to run continuously.

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

- A. Host each website by using AWS Amplify with a serverless backend. Connect the repository branches that correspond to each of the desired environments. Start deployments by merging code changes to a desired branch.
- B. Host each website in AWS Elastic Beanstalk with multiple environments. Use the EB CLI to link each repository branch. Integrate AWS CodePipeline to automate deployments from version control code merges.
- C. Host each website in different Amazon S3 buckets for each environment. Configure AWS CodePipeline to pull source code from version control. Add an AWS CodeBuild stage to copy source code to Amazon S3.
- D. Host each website on its own Amazon EC2 instance. Write a custom deployment script to bundle each website's static assets. Copy the assets to Amazon EC2. Set up a workflow to run the script when code is merged.

Correct Answer: A

Community vote distribution

A (89%) 11%

 **Bibay** Highly Voted 1 year, 8 months ago

a


The solution that will meet these requirements with the LEAST operational overhead is option A: Host each website by using AWS Amplify with a serverless backend. AWS Amplify is a fully managed service that allows developers to build and deploy web applications and static websites. With Amplify, developers can easily connect their repositories, such as AWS CodeCommit, Bitbucket, and GitHub, to automatically build and deploy changes to the website based on code merges. Amplify also supports phased releases with multiple environments, including development, staging, user acceptance testing, and production, which can be linked to specific branches in the repository. Additionally, Amplify uses HTTPS for all data exchange by default and has a serverless backend, which means there are no servers to maintain. Overall, this solution provides the least operational overhead while meeting all the specified requirements.

upvoted 28 times

 **yashika2005** 1 year, 8 months ago

thanks a ton for all the explanations!

upvoted 3 times

 **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: A

The correct answer is A.


AWS Amplify is an all in one service for the requirement.

<https://docs.aws.amazon.com/amplify/latest/userguide/welcome.html>

Option C is almost correct, but it does not mention how to implement HTTPS.

Option B and D are wrong. They need to keep running servers.

upvoted 21 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B

A) AWS Amplify is purpose-built for deploying and hosting static websites and serverless backends with minimal operational overhead.

B) Eliminated - Elastic Beanstalk is overkill for hosting static websites since it is designed for running full-stack web applications with servers

C) Eliminated : S3 does not natively support branch-to-environment linking

D) Eliminated: Using EC2 instances is entirely unnecessary for hosting static websites. It involves the highest operational overhead

upvoted 1 times

🗨️ 👤 **sumanshu** 1 month, 1 week ago
Sorry - By mistake Select "B", should be A
upvoted 1 times

🗨️ 👤 **sumanshu** 1 day, 7 hours ago
C) Eliminated - Additional steps needed to configure CloudFront for HTTPS. Also need multiple S3 buckets (for dev, stg, prod etc).

Though S3 is great choice, but as compare to Amplify, more effort
upvoted 1 times

🗨️ 👤 **trieudo** 1 month, 2 weeks ago

Selected Answer: A

keyword:

- static websites
- version control systems, including AWS CodeCommit, Bitbucket, and GitHub.
- phased releases by using development, staging, user acceptance testing, and production environments
- use HTTPS for all data exchange
- not require servers to run continuously.
- LEAST operational overhead

=> discard D: EC2 instance is not serverless, violate 'not require servers to run continuously'

=> discard C: you must set up many thing manually, create each bucket for each stage,... .Violate 'LEAST operational overhead'

=> Discard B: it excess capacity for 'static websites', usually using when deploy dynamic complexity backend, you must manage EC2, ELB, Auto Scaling, ... so violate 'LEAST operational overhead'

A fits all requirement above

upvoted 1 times

🗨️ 👤 **Saurabh04** 5 months, 2 weeks ago

Selected Answer: A

Amplify hosting provides a git based workflow for hosting full stack server less applications with continuous deployment.

upvoted 1 times

🗨️ 👤 **65703c1** 8 months, 2 weeks ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ 👤 **Vaibs099** 9 months, 1 week ago

Amplify is the best option here to host website / static site as well with Hosting Environment option which can pull code from github, codecommit and bitbucket. Webapp Hosting can be for different envs like Prod, Dev etc. This gives serverless hosting option along with HTTPS. S3 static website hosting has no HTTPS and EB and EC2 are running instances.

upvoted 2 times

🗨️ 👤 **Cerakoted** 1 year, 3 months ago

Selected Answer: A

Check About AWS Amplify Hosting

upvoted 2 times

🗨️ 👤 **jayvarma** 1 year, 5 months ago

Option A is the answer. Ofcourse, until now we have been used to the fact that we need to use S3 for static website hosting.

But there are a lot of requirements described in the question like the source code hosting, phased releases with different environments and HTTPS for all data exchange (which is not possible with S3 Hosting).

AWS Amplify does all of this for you with the least operational overhead.

upvoted 5 times

🗨️ 👤 **Devon_Fazekas** 1 year, 8 months ago

For fellow ACloudGurus, I was taught to associate static website hosting to S3 buckets. But apparently, "least operational overhead" is achieved using Amplify, as it natively supports deployment to various environments and seamlessly integrates with version control systems. Whereas, S3 requires configuring multiple buckets, configuring CodePipeline and integrating with each bucket.

upvoted 4 times

🗨️ 👤 **Rpod** 1 year, 9 months ago

Selected Answer: C

Static Website should be C ..using S3

upvoted 3 times

🗨️ 👤 **Arnaud92** 1 year, 8 months ago

Sadly Static Web Hosting on S3 does not supports HTTPS . So Response is A ;-)

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/WebsiteHosting.html>

upvoted 8 times

🗨️ 👤 **jipark** 1 year, 6 months ago

that is critical key !! thanks a lot.

upvoted 2 times

A company is migrating an on-premises database to Amazon RDS for MySQL. The company has read-heavy workloads. The company wants to refactor the code to achieve optimum read performance for queries.

Which solution will meet this requirement with LEAST current and future effort?

- A. Use a multi-AZ Amazon RDS deployment. Increase the number of connections that the code makes to the database or increase the connection pool size if a connection pool is in use.
- B. Use a multi-AZ Amazon RDS deployment. Modify the code so that queries access the secondary RDS instance.
- C. Deploy Amazon RDS with one or more read replicas. Modify the application code so that queries use the URL for the read replicas.
- D. Use open source replication software to create a copy of the MySQL database on an Amazon EC2 instance. Modify the application code so that queries use the IP address of the EC2 instance.

Correct Answer: C

Community vote distribution

C (97%)

🗳️ **Devon_Fazekas** Highly Voted 1 year, 8 months ago

Selected Answer: C

Multi-AZ is for disaster recovery, not read scalability or performance.

upvoted 7 times

🗳️ **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: C

A) & B) - Eliminated - Multi-AZ deployments are designed for high availability and failover, not for scaling read performance.

D) Eliminated - significant operational overhead to set up and maintain replication between the primary database and the EC2 instance

upvoted 2 times

🗳️ **trieudo** 1 month, 2 weeks ago

Selected Answer: C

keyword: read-heavy workloads, LEAST current and future effort

==> discard A, B: multiAZ, just only make high availability, option A: it make you can handler bigger cocurrency request, option B you can access to secondary RDS in normal case by modifying code. Both are not help for 'read-heavy workloads'

==> discard D: take time to execute, maintain, ... when use not-intergated source, violate 'LEAST current and future effort'

C is best choice

upvoted 2 times

🗳️ **Piku2** 1 month, 2 weeks ago

Selected Answer: C

It should be C as we need to update the URL of the rds endpoint as it is needed to connect the application to use the read replicas for read queries.

upvoted 1 times

🗳️ **tomchandler077** 6 months, 3 weeks ago

Option C provides the most straightforward and effective solution for improving read performance with minimal changes to the current application code and the least ongoing maintenance effort. Deploying read replicas allows for scaling read capacity and distributing read traffic efficiently.

upvoted 3 times

🗳️ **nkroker** 7 months ago

Option C is wrong because deploying a read replica will be more effort then just enabling the multi-AZ with RDS and also the multi-AZ is meant for high availability that's why option B is correct.

upvoted 1 times

🗳️ **JonasKahnwald** 2 months, 2 weeks ago

No, multi AZ is for disaster recovery.

upvoted 1 times

🗳️ 👤 **65703c1** 8 months, 2 weeks ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗳️ 👤 **Vaibs099** 9 months, 1 week ago

Option A and B are both talking about Multi AZ RDS instance which gives Primary and Secondary(Non Read Replica). This is good for high availability but will not help in reads. Read replica or Multi AZ Cluster deployment is the only option to achieve high reads.

upvoted 1 times

🗳️ 👤 **Dikshika** 9 months, 2 weeks ago

C it is as it clearly mentions they want to achieve optimum read performance

upvoted 2 times

🗳️ 👤 **dan_bj** 9 months, 3 weeks ago

Selected Answer: C

C forever

upvoted 1 times

🗳️ 👤 **badsati** 9 months, 3 weeks ago

Selected Answer: C

C... No Question

upvoted 1 times

🗳️ 👤 **vinfo** 9 months, 4 weeks ago

Selected Answer: C

C. El uso de replicas de lectura, aliviana las consultas intensivas sobre la BD principal

upvoted 1 times

🗳️ 👤 **xdkonorek2** 1 year, 1 month ago

Selected Answer: B

easiest solution is to use multi-az rds deployment with 2 readable standby instances

setting up read replica is more effort than checking a single option

upvoted 2 times

🗳️ 👤 **mghectorenjoyer69** 10 months, 1 week ago

ni mada ra

upvoted 1 times

🗳️ 👤 **Skywalker23** 1 year, 4 months ago

Selected Answer: C

Read heavy access need read replicas as the right solution.

upvoted 4 times

🗳️ 👤 **Tony88** 1 year, 5 months ago

Selected Answer: C

Keyword: heavy read

upvoted 3 times

🗳️ 👤 **Akash619** 1 year, 5 months ago

Selected Answer: C

Read Replicas for high performance read operations

upvoted 2 times

🗳️ 👤 **jayvarma** 1 year, 5 months ago

Keyword: Achieve Optimum read performance for queries.

Answer: Use Read Replicas and use that specific URL for read queries.

upvoted 2 times

A developer is creating an application that will be deployed on IoT devices. The application will send data to a RESTful API that is deployed as an AWS Lambda function. The application will assign each API request a unique identifier. The volume of API requests from the application can randomly increase at any given time of day.

During periods of request throttling, the application might need to retry requests. The API must be able to handle duplicate requests without inconsistencies or data loss.

Which solution will meet these requirements?

- A. Create an Amazon RDS for MySQL DB instance. Store the unique identifier for each request in a database table. Modify the Lambda function to check the table for the identifier before processing the request.
- B. Create an Amazon DynamoDB table. Store the unique identifier for each request in the table. Modify the Lambda function to check the table for the identifier before processing the request.
- C. Create an Amazon DynamoDB table. Store the unique identifier for each request in the table. Modify the Lambda function to return a client error response when the function receives a duplicate request.
- D. Create an Amazon ElastiCache for Memcached instance. Store the unique identifier for each request in the cache. Modify the Lambda function to check the cache for the identifier before processing the request.

Correct Answer: B

Community vote distribution

B (100%)

🗳️ **Devon_Fazekas** Highly Voted 1 year, 8 months ago

Selected Answer: B

I originally thought ElastiCache would provide the sufficient session management of the unique identifiers with the least latency. But apparently, the scope of this question revolves around durability, not latency. Hence, a persistent storage is better suited. And while RDS is a viable solution for durability and performance, the question specifies IoT devices which typically produce unstructured data that is better handled by No-SQL services like DynamoDB.

upvoted 29 times

🗳️ **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: B

B

The resolution is to make the Lambda function idempotent.

<https://repost.aws/knowledge-center/lambda-function-idempotent>

<https://aws.amazon.com/builders-library/making-retries-safe-with-idempotent-APIs/>

upvoted 9 times

🗳️ **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B

A) Eliminated - RDS is not as scalable as DynamoDB for handling unpredictable spikes in traffic.

D) Eliminated - Memcached is an in-memory cache that is not designed for persistence or durability. If the instance is restarted or the data is evicted due to memory constraints, the identifiers could be lost

upvoted 3 times

🗳️ **sumanshu** 1 month, 1 week ago

C) Eliminated - Returning a client error response for duplicate requests is not a good approach

upvoted 2 times

🗳️ **trieudo** 1 month, 2 weeks ago

Selected Answer: B

keyword:

- AWS Lambda function
- increase at any given time
- retry requests
- handle duplicate requests

=> Discard A, D: violate 'increase at any given time', it is not automatically, because it is not fit with serverless service 'AWS Lambda function'
=> Discard C: violate 'retry requests + handle duplicate requests', => different response from success vs. error error, make it is non-idempotency
=> Discard D: violate 'inconsistencies or data loss.' it does not store persistent data.

B is best solution
upvoted 3 times

🗨️ **f271c23** 1 month, 3 weeks ago

Selected Answer: C

I think C has the merits to be a right answer. It has more specific handling to duplicate requests . The answer B does the same but the language is more explicit in option C
upvoted 1 times

🗨️ **ElFaramawi** 8 months ago

Why not C?

According to the question "During periods of request throttling, the application might need to retry requests", this indicate that lambda should returns client error, so the application can make another retry request to fix the problem
upvoted 3 times

🗨️ **queekao** 6 months, 2 weeks ago

because c don't mention about retry actions
upvoted 1 times

🗨️ **cocolavayen** 7 months ago

same doubt
upvoted 1 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: B

B is the correct answer.
upvoted 1 times

🗨️ **Abdullah22** 10 months, 3 weeks ago

Selected Answer: B

Consistency: Memcached does not provide built-in support for atomic operations or conditional writes like DynamoDB does. Handling duplicate requests and ensuring consistency would require additional application logic and complexity.
upvoted 1 times

🗨️ **Tony88** 1 year, 5 months ago

Selected Answer: B

Cache topic.

So Elastic Redis and DynamoDB both can be used as a cache solution.

If you want high performance, low latency, go with Redis

If you want persistent storage, go with DyanmoDB.

upvoted 5 times

A developer wants to expand an application to run in multiple AWS Regions. The developer wants to copy Amazon Machine Images (AMIs) with the latest changes and create a new application stack in the destination Region. According to company requirements, all AMIs must be encrypted in all Regions. However, not all the AMIs that the company uses are encrypted.

How can the developer expand the application to run in the destination Region while meeting the encryption requirement?

- A. Create new AMIs, and specify encryption parameters. Copy the encrypted AMIs to the destination Region. Delete the unencrypted AMIs.
- B. Use AWS Key Management Service (AWS KMS) to enable encryption on the unencrypted AMIs. Copy the encrypted AMIs to the destination Region.
- C. Use AWS Certificate Manager (ACM) to enable encryption on the unencrypted AMIs. Copy the encrypted AMIs to the destination Region.
- D. Copy the unencrypted AMIs to the destination Region. Enable encryption by default in the destination Region.

Correct Answer: A

Community vote distribution

A (75%)

B (25%)

 **Bibay** Highly Voted 1 year, 8 months ago

A. Create new AMIs, and specify encryption parameters. Copy the encrypted AMIs to the destination Region. Delete the unencrypted AMIs.

The best solution for meeting the encryption requirement is to create new AMIs with encryption enabled and copy them to the destination Region. By default, when an AMI is copied to another Region, it is not encrypted in the destination Region even if it is encrypted in the source Region. Therefore, the developer must create new encrypted AMIs that can be used in the destination Region. Once the new encrypted AMIs have been created, they can be copied to the destination Region. The unencrypted AMIs can then be deleted to ensure that all instances running in all Regions are using only encrypted AMIs.

upvoted 25 times


 **Rameez1** Highly Voted 1 year, 3 months ago

Selected Answer: A

A is correct.

Unencrypted AMI can't be encrypted after creation. Need to create new encrypted AMI then it can be copied to other regions.

upvoted 10 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: A

A) Amazon Machine Images (AMIs) can be encrypted at creation time

B) Eliminated - AWS KMS does not allow you to encrypt an existing unencrypted AMI directly.

C) Eliminated - AWS Certificate Manager (ACM) is used for managing SSL/TLS certificates, not for encrypting AMIs.


D) Eliminated - It does not retroactively encrypt existing AMIs.

upvoted 2 times

 **sumanshu** 1 day, 7 hours ago

D) Eliminated - Though we have the option while copying the AMI to different region, we can encrypt at that time. But it will only encrypt in destination region. What about the Source region. As per question - It should be encrypted in all regions.

upvoted 1 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: A

==> Discard B: Once an AMI is created, encryption configuration cannot be changed,

==> Discard C: ACM use for SSL/ TLS connection manager

==> Discard D: Even if assumed that "encryption by default" is enabled in the destination before copy, original AMI is still not encrypted, so condition "AMIs must be encrypted in all Regions" is not met.

A is popular pattern:

1. Create a snapshot from the original AMI.
2. Encrypt the snapshot using an AWS KMS key.
3. Create a new AMI from the encrypted snapshot.
4. Copy the encrypted AMI to the destination region

upvoted 1 times

🗨️ 👤 **Venky786** 2 months, 3 weeks ago

Answer is A

While AWS KMS is used to manage encryption keys, it cannot retroactively encrypt an existing unencrypted AMI. Encryption must be specified when creating or copying the AMI.

upvoted 1 times

🗨️ 👤 **wh1t4k3r** 5 months, 2 weeks ago

Selected Answer: A

Regarding B: Once an AMI is created, encryption configuration cannot be changed, you need to create a new one and enable encryption. Another point: if you are planning to share the AMI between accounts, you cannot use AWS managed keys

upvoted 1 times

🗨️ 👤 **tomchandler077** 6 months, 3 weeks ago

Option A ensures all AMIs are encrypted before they are copied to the destination region, meeting the encryption requirement and providing a clear and compliant process for expanding the application to multiple AWS Regions.

upvoted 1 times

🗨️ 👤 **65703c1** 8 months, 2 weeks ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ 👤 **TheFivePips** 11 months ago

Selected Answer: A

Encryption of an Amazon Machine Image (AMI) is typically tied to the underlying Amazon Elastic Block Store (EBS) snapshots that are associated with the AMI.

When you create an AMI, you have the option to specify encryption parameters. If you choose to encrypt the root volume, the resulting AMI will be encrypted. This encryption setting applies to both the root volume and any additional EBS volumes attached to the instance.

The encryption status of an EBS snapshot is determined at the time of snapshot creation. Once a snapshot is created, its encryption status remains constant. If you want to encrypt a snapshot, you typically need to create a new snapshot from an encrypted volume.

Once an AMI is created, you generally cannot modify its encryption status directly. If you need to change the encryption status, you might need to create a new AMI from an encrypted snapshot.

upvoted 4 times

🗨️ 👤 **SerialiDr** 11 months, 1 week ago

Selected Answer: A

A. This approach ensures that all AMIs are encrypted using specified encryption parameters before they are copied to the destination Region, aligning with the company's encryption requirement. AWS provides the capability to encrypt AMIs during the AMI creation process and when copying AMIs between Regions. You can specify an AWS Key Management Service (AWS KMS) customer master key (CMK) during these processes to use for encryption, meeting the requirement to use a company-generated key.

upvoted 1 times

🗨️ 👤 **gqs3119** 1 year, 1 month ago

C ACM is about SSL/TLS

D Even if assumed that "encryption by default" is enabled in the destination before copy, original AMI is still not encrypted, so condition "AMIs must be encrypted in all Regions" is not met.

B I don't see any option in AWS Console or docs to encrypt in place existing AMI. It can be done when copying it. Option B doesn't handle existing unencrypted AMIs.

A I think, A is the best description of the procedure.

upvoted 2 times

🗨️ 👤 **BluntFarmer** 1 year, 2 months ago

I would go with D: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html#encryption-by-default>

Solves must be encrypted issue once and for all plus you can copy unencrypted to encrypted

upvoted 2 times

🗨️ **maurice2005** 11 months, 3 weeks ago

it still keeps the unencrypted AMI untouched. You have to delete them but not mentioned as explicit as A

upvoted 2 times

🗨️ **walala97** 1 year, 2 months ago

Selected Answer: A

kms keys is regional,so when you use kms before you copy to another region,the second region still has the unencryed AMIs.so B is not correct

upvoted 1 times

🗨️ **ronn555** 1 year, 2 months ago

A

When you create an encrypted AMI and do not specify the KMS key, AWS will use the default Customer Managed Key which is the only multi-region key. If you select a KMS key from the origin region it will not work in the destination region (presently) so B is not correct.

upvoted 3 times

🗨️ **Cerakoted** 1 year, 3 months ago

Selected Answer: B

Answer is B

check this link

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/CopyingAMIs.html#ami-copy-encryption>

upvoted 2 times

🗨️ **[Removed]** 1 year, 1 month ago

If you read this link carefully it actually proves that B is wrong. The correct answer is A. You cannot enable encryption on an unencrypted AMI. ---> an AMI backed by an unencrypted root snapshot is copied to an AMI with an encrypted root snapshot. The CopyImage action is invoked with two encryption parameters, including a customer managed key. As a result, the encryption status of the root snapshot changes, so that the target AMI is backed by a root snapshot containing the same data as the source snapshot, but encrypted using the specified key.

upvoted 1 times

🗨️ **manikantaJ** 1 year, 3 months ago

Selected Answer: B

Here's why option B is the appropriate choice:

AWS KMS Encryption: AWS KMS is a service that allows you to easily enable encryption for your resources, including Amazon Machine Images (AMIs). You can create a customer managed key (CMK) in AWS KMS and use it to encrypt your AMIs.

Enable Encryption on Unencrypted AMIs: You can enable encryption for unencrypted AMIs by creating a copy of the AMI and specifying the AWS KMS key to use for encryption during the copy process. This ensures that your new AMIs in the destination Region are encrypted.

Maintain Data Integrity: This approach allows you to maintain data integrity and ensure that all AMIs are encrypted in compliance with company requirements.

upvoted 2 times

🗨️ **sofiatian** 1 year, 4 months ago

Selected Answer: B

Copy an unencrypted source AMI to an encrypted target AMI

In this scenario, an AMI backed by an unencrypted root snapshot is copied to an AMI with an encrypted root snapshot. The CopyImage action is invoked with two encryption parameters, including a customer managed key. As a result, the encryption status of the root snapshot changes, so that the target AMI is backed by a root snapshot containing the same data as the source snapshot, but encrypted using the specified key.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/CopyingAMIs.html>

upvoted 2 times

A company hosts a client-side web application for one of its subsidiaries on Amazon S3. The web application can be accessed through Amazon CloudFront from <https://www.example.com>. After a successful rollout, the company wants to host three more client-side web applications for its remaining subsidiaries on three separate S3 buckets.

To achieve this goal, a developer moves all the common JavaScript files and web fonts to a central S3 bucket that serves the web applications. However, during testing, the developer notices that the browser blocks the JavaScript files and web fonts.


What should the developer do to prevent the browser from blocking the JavaScript files and web fonts?

- A. Create four access points that allow access to the central S3 bucket. Assign an access point to each web application bucket.
- B. Create a bucket policy that allows access to the central S3 bucket. Attach the bucket policy to the central S3 bucket
- C. Create a cross-origin resource sharing (CORS) configuration that allows access to the central S3 bucket. Add the CORS configuration to the central S3 bucket.
- D. Create a Content-MD5 header that provides a message integrity check for the central S3 bucket. Insert the Content-MD5 header for each web application request.

Correct Answer: C

Community vote distribution

C (100%)

 **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: C

C

This is a frequent trouble. Web applications cannot access the resources in other domains by default, except some exceptions. You must configure CORS on the resources to be accessed.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/cors.html>

upvoted 6 times


 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: C

The issue described is related to Cross-Origin Resource Sharing (CORS).

- A) Eliminated - S3 access points deals with permissions, not with resolving cross-origin issues.
- B) Eliminated - A bucket policy defines access rules for an S3 bucket (e.g., which users or accounts can access it). Bucket policies define access permissions, but they do not address the cross-origin issue.
- C) By adding a CORS configuration to the central bucket, you instruct the browser to allow requests from the web application's domains (origins).
- D) Eliminated - Content-MD5 is a header used to ensure that data was not corrupted during transmission. it is unrelated to the browser's blocking behavior due to CORS.

upvoted 2 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: C

==> Discard A. Access points manage S3 access but don't address browser's cross-origin restrictions.

==> Discard B. Bucket policies control permissions but don't override browser Same-Origin Policy.

==> Discard D. Content-MD5 ensures data integrity but doesn't affect cross-origin resource sharing.

C is useful when cross site access

upvoted 1 times

 **tomchandler077** 6 months, 3 weeks ago

C

The question described is a classic case of Cross-Origin Resource Sharing (CORS) where the browser blocks resources (like JavaScript files and web fonts) that are loaded from a different origin (the central S3 bucket) than the web application. To resolve this, a CORS configuration needs to be added to the central S3 bucket to allow these resources to be accessed from the different origins of the web applications.

upvoted 1 times

🗨️ 👤 **65703c1** 8 months, 2 weeks ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗨️ 👤 **badsati** 9 months, 3 weeks ago

Selected Answer: C

The answer is C for cat

upvoted 1 times

🗨️ 👤 **svrvtr** 1 year, 10 months ago

Selected Answer: C

It is C

upvoted 3 times

🗨️ 👤 **aragon_saa** 1 year, 10 months ago

C

<https://www.examttopics.com/discussions/amazon/view/88856-exam-aws-certified-developer-associate-topic-1-question-302/>

upvoted 3 times

An application is processing clickstream data using Amazon Kinesis. The clickstream data feed into Kinesis experiences periodic spikes. The PutRecords API call occasionally fails and the logs show that the failed call returns the response shown below:

```
{
  "FailedRecordCount": 1,
  "Records": [
    {
      "SequenceNumber": "21269319989900637946712965403778482371",
      "ShardId": "shardId-000000000001"
    },
    {
      "ErrorCode": "ProvisionedThroughputExceededException",
      "ErrorMessage": "Rate exceeded for shard shardId-000000000001 in
        stream exampleStreamName under account 123456789."
    },
    {
      "SequenceNumber": "21269319989999637946712965403778482985",
      "ShardId": "shardId-000000000002"
    }
  ]
}
```

Which techniques will help mitigate this exception? (Choose two.)

- A. Implement retries with exponential backoff.
- B. Use a PutRecord API instead of PutRecords.
- C. Reduce the frequency and/or size of the requests.
- D. Use Amazon SNS instead of Kinesis.
- E. Reduce the number of KCL consumers.

Correct Answer: AC

Community vote distribution

AC (82%)

BC (18%)

 **eboehm2** Highly Voted 1 year, 7 months ago


Selected Answer: AC

100% AC as per AWS : ProvisionedThroughputExceededException

The request rate for the stream is too high, or the requested data is too large for the available throughput. Reduce the frequency or size of your requests. For more information, see Streams Limits in the Amazon Kinesis Data Streams Developer Guide, and Error Retries and Exponential Backoff in AWS in the AWS General Reference.

https://docs.aws.amazon.com/kinesis/latest/APIReference/API_PutRecords.html

upvoted 8 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: AC

A) Correct - Retries with exponential backoff help manage temporary spikes by delaying subsequent retry attempts progressively

B) Eliminated - Switching to PutRecord would increase API calls, worsening the throughput issue

C) Correct - Reducing request size or frequency directly addresses the throughput limits by spreading out the data more evenly.

D) Eliminated - Amazon SNS is a messaging service, not designed for streaming data with high throughput like Kinesis.

upvoted 1 times

 **sumanshu** 1 month, 1 week ago

E) Eliminated - KCL consumers consume data from Kinesis streams; reducing their number will not address the write throughput limits causing the ProvisionedThroughputExceededException

upvoted 1 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: AC

==> Discard B: PutRecord doesn't reduce throughput issues as it sends one record at a time, increasing API calls.

==> Discard D: SNS isn't designed for streaming data; it handles pub/sub messaging.

==> Discard E: Reducing KCL consumers affects reading, not writing throughput.

- A: Exponential backoff reduces retries during peak usage, easing shard throughput pressure.

- C: Lowering request frequency or size directly mitigates throughput exceedance on shards.

upvoted 1 times

🗨️ **Venky786** 2 months, 3 weeks ago

Answer is A and C

Batch Efficiency: The `PutRecords` API allows you to send multiple records in a single request, which is generally more efficient than sending individual records with `PutRecord`. Using `PutRecords` can help optimize throughput by reducing the number of API calls and better utilizing the available capacity.

upvoted 2 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: AC

AC is the correct answer.

upvoted 1 times

🗨️ **Baba_Eni** 1 year, 8 months ago

Selected Answer: AC

AC is the best answer. When there is throttling, it is best practise to implement retries with exponential backoff.

upvoted 1 times

🗨️ **ezredame** 1 year, 8 months ago

Selected Answer: BC

I think this is really tricky question. To get this exception, the request rate for the stream is too high, or the requested data is too large for the available throughput. Reduce the frequency or size of your requests. So we can "Reduce the frequency and/or size of the requests" also decrease the size with "Use a PutRecord API instead of PutRecords"

The API already implements retries with exponential backoff. So there is no need for A.

upvoted 3 times

🗨️ **eboehm2** 1 year, 7 months ago

I thought this at first too, but I was doing some additional reading and using the PutRecord API over PutRecords is wrong as it could actually make the problem worse as producers may make too many rapid requests to write to the stream

<https://repost.aws/knowledge-center/kinesis-data-stream-throttling>

upvoted 4 times

🗨️ **Majong** 1 year, 8 months ago

Can you please add a link where I can find this information. From what I can read on AWS is that you can implement exponential backoff but it is not by default.

upvoted 1 times

🗨️ **Untamables** 1 year, 10 months ago

Selected Answer: AC

A and C

<https://aws.amazon.com/premiumsupport/knowledge-center/kinesis-data-stream-throttling-errors/>

upvoted 4 times

🗨️ **aragon_saa** 1 year, 10 months ago

AC

<https://www.examttopics.com/discussions/amazon/view/69142-exam-aws-certified-developer-associate-topic-1-question-370/>

upvoted 4 times

🗨️ **yashika2005** 1 year, 8 months ago

thanks a lotttt!

upvoted 1 times

A company has an application that uses Amazon Cognito user pools as an identity provider. The company must secure access to user records. The company has set up multi-factor authentication (MFA). The company also wants to send a login activity notification by email every time a user logs in.

What is the MOST operationally efficient solution that meets this requirement?

- A. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notification. Add an Amazon API Gateway API to invoke the function. Call the API from the client side when login confirmation is received.
- B. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notification. Add an Amazon Cognito post authentication Lambda trigger for the function.
- C. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notification. Create an Amazon CloudWatch Logs log subscription filter to invoke the function based on the login status.
- D. Configure Amazon Cognito to stream all logs to Amazon Kinesis Data Firehose. Create an AWS Lambda function to process the streamed logs and to send the email notification based on the login status of each user.

Correct Answer: B

Community vote distribution

B (91%)

9%


 **Bibay** Highly Voted 1 year, 8 months ago

B. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notification. Add an Amazon Cognito post authentication Lambda trigger for the function.

The most operationally efficient solution for sending login activity notifications by email for Amazon Cognito user pools is to use a Lambda trigger that is automatically invoked by Amazon Cognito every time a user logs in. This eliminates the need for client-side calls to an API or log subscription filter. A Lambda function can be used to send email notifications using Amazon SES.

Option B satisfies these requirements and is the most operationally efficient solution.

upvoted 13 times


 **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: B

B

<https://docs.aws.amazon.com/cognito/latest/developerguide/user-pool-lambda-post-authentication.html>

upvoted 9 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B


A) Eliminated - the API Gateway needs to be explicitly invoked from the client. not operationally efficient

B) Correct - Post-authentication triggers are built-in Cognito features that automatically execute a Lambda function after a user logs in.

C) Eliminated - less efficient because it requires setting up CloudWatch Logs for Cognito, creating subscription filters, and maintaining log processing infrastructure

D) Eliminated - highly complex and introduces multiple components (Kinesis, Firehose, Lambda),

upvoted 1 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: B

==> Discard A: Requires client-side API calls, which are less secure and operationally inefficient (enhance security, validate in BE, ...) compared to server-side triggers.

==> Discard C: Relies on CloudWatch Logs and filters, which add unnecessary complexity and are not directly tied to Cognito's login events.

==> Discard D: Involves Kinesis Data Firehose for log streaming, which is over-engineered and introduces additional cost and latency for a simple notification task.

B: Uses Cognito's native post-authentication trigger, which is the most secure, integrated, and operationally efficient method for sending email notifications after successful logins.

upvoted 1 times

🗨️ **serverlessme** 3 months, 1 week ago

Option B is absolutely correct.

The Lambda function is triggered by Cognito whenever a user log in. The Lambda function then sends email notification to the user using Amazon SES.

upvoted 1 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗨️ **Melisa202401** 9 months, 4 weeks ago

Selected Answer: A

Amazon Cognito user pools integrate with API Gateway or ALB

Process is: user authenticate with CUP, receive JWT (token), then pass to API Gateway

API Gateway will evaluate JWT with CUP, if it is valid, allow access to Lambda (have a duty to send email)

upvoted 1 times

🗨️ **Melisa202401** 9 months, 3 weeks ago

sorry I change to B

Because the question have the presence of MFA

upvoted 1 times

🗨️ **aragon_saa** 1 year, 10 months ago

B

<https://www.examttopics.com/discussions/amazon/view/78944-exam-aws-certified-developer-associate-topic-1-question-9/>

upvoted 3 times

A developer has an application that stores data in an Amazon S3 bucket. The application uses an HTTP API to store and retrieve objects. When the PutObject API operation adds objects to the S3 bucket the developer must encrypt these objects at rest by using server-side encryption with Amazon S3 managed keys (SSE-S3).

Which solution will meet this requirement?

- A. Create an AWS Key Management Service (AWS KMS) key. Assign the KMS key to the S3 bucket.
- B. Set the x-amz-server-side-encryption header when invoking the PutObject API operation.
- C. Provide the encryption key in the HTTP header of every request.
- D. Apply TLS to encrypt the traffic to the S3 bucket.

Correct Answer: B

Community vote distribution



Bibay Highly Voted 1 year, 2 months ago

B. Set the x-amz-server-side-encryption header when invoking the PutObject API operation.

When using the PutObject API operation to store objects in an S3 bucket, the x-amz-server-side-encryption header can be set to specify the server-side encryption algorithm used to encrypt the object. Setting this header to "AES256" or "aws:kms" enables server-side encryption with SSE-S3 or SSE-KMS respectively.

Option A is incorrect because assigning a KMS key to the S3 bucket will not enable SSE-S3 encryption.

Option C is incorrect because providing the encryption key in the HTTP header of every request is not a valid way to enable SSE-S3 encryption.

Option D is incorrect because applying TLS encryption to the traffic to the S3 bucket only encrypts the data in transit, but does not encrypt the objects at rest in the bucket.

upvoted 15 times

jipark 12 months ago

I now got to know 'KMS key to S3 bucket will not enable S3 encryption'

upvoted 1 times

beekeeper0101 1 month, 3 weeks ago

Thank you!

Setting the x-amz-server-side-encryption header to:

- AES256 => SSE-S3

- AWS:KMS => SSE-KMS

upvoted 1 times

svrnvtr Highly Voted 1 year, 4 months ago

Selected Answer: B

B <https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingServerSideEncryption.html>

upvoted 10 times

sumanshu Most Recent 1 month, 1 week ago

Selected Answer: B

A) Eliminated - While AWS Key Management Service (KMS) keys can be used for encryption in S3 (SSE-KMS), this option refers to creating a custom KMS key, which is not required when using SSE-S3.

B) Correct - The x-amz-server-side-encryption header is the correct way to specify the use of SSE-S3 when uploading objects to S3 via the PutObject API.

C) Eliminated - Providing an encryption key in the HTTP header refers to client-side encryption or SSE-C

D) Eliminated - TLS (Transport Layer Security) encrypts data in transit, not at rest.

upvoted 1 times

trieudo 1 month, 2 weeks ago



Selected Answer: B

==> Discard A: SSE-KMS uses AWS KMS keys, not Amazon S3-managed keys required for SSE-S3.

==> Discard C: SSE-C requires customer-provided keys, not Amazon S3-managed keys for SSE-S3.

==> Discard D: TLS encrypts data in transit, not at rest as required by SSE-S3.


B is correct because setting `x-amz-server-side-encryption: AES256` ensures Amazon S3 uses SSE-S3 to encrypt objects at rest automatically.
upvoted 1 times

  **65703c1** 2 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

  **badsati** 3 months, 3 weeks ago

Selected Answer: B

Answer is B

upvoted 1 times

  **nderitunick** 8 months ago

Aren't objects on s3 encrypted using SSE-S3 by default? I don't understand why D is not the answer.

upvoted 1 times

  **nderitunick** 8 months ago

I misread the question. It's all good.

upvoted 1 times

  **aanataliya** 11 months, 1 week ago

Answer for this question is changed starting January 5, 2023. Amazon S3 now applies server-side encryption with Amazon S3 managed keys (SSE-S3) as the base level of encryption for every bucket in Amazon S3.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/default-encryption-faq.html>

upvoted 8 times

  **fordiscussiontwo** 10 months ago

what is correct answer then?

upvoted 2 times

  **cucuff** 7 months, 2 weeks ago

because it takes some time for exam questions to be updated

upvoted 1 times

  **[Removed]** 1 year ago

Selected Answer: B

Header parameter "s3:x-amz-server-side-encryption": "AES256"

upvoted 4 times

  **ttamttam** 1 year ago

Selected Answer: B

C is a way to use customer-provided keys not S3-managed keys.

upvoted 2 times

  **CisconAWSGURU** 1 year, 1 month ago

Selected Answer: C

C is correct and here is the reason from AWS docs.

Visit AWS Regions and Endpoints in the AWS General Reference or the AWS Region Table to see the regional availability for ACM.

Certificates in ACM are regional resources. To use a certificate with Elastic Load Balancing for the same fully qualified domain name (FQDN) or set of FQDNs in more than one AWS region, you must request or import a certificate for each region. For certificates provided by ACM, this means you must revalidate each domain name in the certificate for each region. You cannot copy a certificate between regions.

To use an ACM certificate with Amazon CloudFront, you must request or import the certificate in the US East (N. Virginia) region. ACM certificates in this region that are associated with a CloudFront distribution are distributed to all the geographic locations configured for that distribution.

upvoted 1 times

A developer needs to perform geographic load testing of an API. The developer must deploy resources to multiple AWS Regions to support the load testing of the API.

How can the developer meet these requirements without additional application code?

- A. Create and deploy an AWS Lambda function in each desired Region. Configure the Lambda function to create a stack from an AWS CloudFormation template in that Region when the function is invoked.
- B. Create an AWS CloudFormation template that defines the load test resources. Use the AWS CLI create-stack-set command to create a stack set in the desired Regions.
- C. Create an AWS Systems Manager document that defines the resources. Use the document to create the resources in the desired Regions.
- D. Create an AWS CloudFormation template that defines the load test resources. Use the AWS CLI deploy command to create a stack from the template in each Region.

Correct Answer: B

Community vote distribution

B (96%) 4%


 **Bibay** Highly Voted 1 year, 8 months ago

Selected Answer: B

B. Create an AWS CloudFormation template that defines the load test resources. Use the AWS CLI create-stack-set command to create a stack set in the desired Regions.

AWS CloudFormation StackSets allow developers to deploy CloudFormation stacks across multiple AWS accounts and regions with a single CloudFormation template. By using the AWS CLI create-stack-set command, the developer can deploy the same CloudFormation stack to multiple regions without additional application code, thereby meeting the requirement for geographic load testing of an API.

upvoted 12 times

 **Untamables** Highly Voted 1 year, 10 months ago


Selected Answer: B

B

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-concepts.html>

<https://awscli.amazonaws.com/v2/documentation/api/2.1.30/reference/cloudformation/create-stack-set.html>

upvoted 8 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B

A) Eliminated: Requires extra steps (deploying Lambda functions) and additional code, making it not optimal.

B) Correct - CloudFormation stack sets are specifically designed for deploying resources across multiple Regions and accounts.

C) Eliminated - AWS Systems Manager documents are used to automate tasks and manage instances, not specifically for deploying infrastructure across multiple Regions.

upvoted 1 times

 **sumanshu** 1 month, 1 week ago

D) Eliminated - Requires you to manually deploy the CloudFormation stack in each Region individually using the deploy command.

upvoted 1 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: B

==> Discard A: Requires additional application code (Lambda), violating "without additional application code.". Beside, Creating and managing multiple Lambda functions increases workload and is inefficient for large-scale deployment.

==> Discard C: AWS Systems Manager is not designed for infrastructure deployment, violating "deploy resources to multiple AWS Regions."

==> Discard D: Requires manual deployment per Region, violating "deploy resources to multiple AWS Regions" efficiently. This means you have to repeat the deployment process manually for each region, which is time-consuming and error-prone.

B: CloudFormation StackSet automates multi-region deployments, meeting "deploy resources to multiple AWS Regions" efficiently without additional application code.

upvoted 1 times

🗨️ 👤 **Itfalcon** 4 months, 2 weeks ago

B according to chatgpt :)

upvoted 1 times

🗨️ 👤 **65703c1** 8 months, 2 weeks ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗨️ 👤 **hsinchang** 1 year, 4 months ago

in desired Regions better than in each Region.

upvoted 3 times

🗨️ 👤 **rlnd2000** 1 year, 5 months ago

Selected Answer: C

If using Edge-Optimized endpoint, then the certificate must be in us-east-1

If using Regional endpoint, the certificate must be in the API Gateway region

upvoted 1 times

🗨️ 👤 **svrvtr** 1 year, 10 months ago

Selected Answer: B

B

<https://aws.amazon.com/ru/about-aws/whats-new/2021/04/deploy-cloudformation-stacks-concurrently-across-multiple-aws-regions-using-aws-cloudformation-stacksets/>

upvoted 3 times

A developer is creating an application that includes an Amazon API Gateway REST API in the us-east-2 Region. The developer wants to use Amazon CloudFront and a custom domain name for the API. The developer has acquired an SSL/TLS certificate for the domain from a third-party provider.

How should the developer configure the custom domain for the application?

- A. Import the SSL/TLS certificate into AWS Certificate Manager (ACM) in the same Region as the API. Create a DNS A record for the custom domain.
- B. Import the SSL/TLS certificate into CloudFront. Create a DNS CNAME record for the custom domain.
- C. Import the SSL/TLS certificate into AWS Certificate Manager (ACM) in the same Region as the API. Create a DNS CNAME record for the custom domain.
- D. Import the SSL/TLS certificate into AWS Certificate Manager (ACM) in the us-east-1 Region. Create a DNS CNAME record for the custom domain.

Correct Answer: D

Community vote distribution

D (86%)

14%

 **brandon87** Highly Voted 1 year, 4 months ago

Selected Answer: D

To use a certificate in AWS Certificate Manager (ACM) to require HTTPS between viewers and CloudFront, make sure you request (or import) the certificate in the US East (N. Virginia) Region (us-east-1).

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/cnames-and-https-requirements.html>

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/CNAMEs.html>

upvoted 31 times

 **ancomedian** Highly Voted 1 year ago

Selected Answer: D

I have checked at various places


Answer is D

Reason: ACM just can only import certificate in us-east-1 and we need to associate the imported certificate with us-east-2

The caused confusion regarding it is because of import and associate

Crux: we will import in us-east-1 but use in us-east-2


upvoted 8 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: D

ACM certificates must reside in us-east-1 for CloudFront, not the same Region as the API (us-east-2 in this case).

upvoted 1 times

 **65703c1** 2 months, 1 week ago

Selected Answer: D

D is the correct answer.


upvoted 1 times

 **fhuadeen** 3 months, 3 weeks ago

Selected Answer: D

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/cnames-and-https-requirements.html>

upvoted 1 times

 **AjeshA1990** 6 months, 3 weeks ago



Import cert in the same region

upvoted 1 times

 **Jonalb** 9 months, 1 week ago

D. Importe o certificado SSL/TLS para o AWS Certificate Manager (ACM) na região us-east-1. Crie um registro DNS CNAME para o domínio personalizado.

upvoted 1 times

  **fossil123** 11 months, 1 week ago

Selected Answer: D

AWS Region for AWS Certificate Manager

To use a certificate in AWS Certificate Manager (ACM) to require HTTPS between viewers and CloudFront, make sure you request (or import) the certificate in the US East (N. Virginia) Region (us-east-1).

upvoted 2 times

  **acordovam** 1 year ago

Selected Answer: D

D

If you need to use CloudFront, then, you must import it into ue-east-1.

<https://docs.aws.amazon.com/acm/latest/userguide/import-certificate.html>

upvoted 3 times

  **Pupina** 1 year ago

Selected Answer: D

A is not right because for cloudfront you create a CNMA not a DNS A

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/CNAMEs.html>

C is not right because ACM cannot import certificates in us-east-2

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/cnames-and-https-requirements.html>

B is not right. The certificate is for an external CA but can be uploaded to ACM or you must request a public certificate from AWS certificate Manager <https://repost.aws/knowledge-center/install-ssl-cloudfront> but you cannot import the certificate into CloudFront

upvoted 3 times



  **rInd2000** 1 year, 1 month ago

Selected Answer: C

C

The first statement of the question: A developer is creating an application that includes an Amazon API Gateway REST API in the us-east-2 Region. ... it is a Regional API, when using a Regional endpoint, the SSL/TLS certificate for the custom domain must be imported into AWS Certificate Manager (ACM) in the same Region as the API, only if we use g Edge-Optimized endpoint, the certificate must be in us-east-1.

upvoted 2 times

  **KarBiswa** 7 months, 2 weeks ago

Initially I also thought but it is a specific hard core requirement "To use an ACM certificate with CloudFront, make sure you request (or import) the certificate in the US East (N. Virginia) Region (us-east-1)."



upvoted 1 times

  **peterpain** 1 year, 2 months ago

Selected Answer: D

The ACM has to be implemented at US-East-1

upvoted 3 times

  **Bibay** 1 year, 2 months ago

Selected Answer: C

To use Amazon CloudFront and a custom domain name for an Amazon API Gateway REST API, the developer should import the SSL/TLS certificate into AWS Certificate Manager (ACM) in the same Region as the API, and create a DNS CNAME record for the custom domain. This is because AWS Certificate Manager can only issue SSL/TLS certificates in the same Region as the API, and a DNS CNAME record maps the custom domain to the CloudFront distribution.

Option A is incorrect because a DNS A record is not sufficient to map the custom domain to the CloudFront distribution.

Option B is incorrect because AWS Certificate Manager must issue the SSL/TLS certificate in the same Region as the API.

Option D is incorrect because the SSL/TLS certificate must be issued in the same Region as the API, and a DNS CNAME record is required to map the custom domain to the CloudFront distribution.

upvoted 5 times

  **KhyatiChhajer** 1 year, 2 months ago

Selected Answer: C

C. Import the SSL/TLS certificate into AWS Certificate Manager (ACM) in the same Region as the API. Create a DNS CNAME record for the custom domain.

Explanation:

Amazon CloudFront can use SSL/TLS certificates stored in AWS Certificate Manager (ACM) to provide secure HTTPS connections for custom domain names. In this scenario, the developer should import the SSL/TLS certificate acquired from a third-party provider into ACM in the same Region as the API (us-east-2 in this case). This allows the certificate to be used by CloudFront.

upvoted 1 times

🗨️ 👤 **hanJR** 1 year, 3 months ago

It's D. It is trying to integrate with CloudFront, therefore it must upload certificates in us-east-1. If it was a regional API, then certificates must be uploaded in the same region of the API Gateway.

upvoted 1 times

🗨️ 👤 **March2023** 1 year, 4 months ago

Selected Answer: C

I was thinking this answer would be C

upvoted 1 times

🗨️ 👤 **Untamables** 1 year, 4 months ago

Selected Answer: D

The correct answer is D.

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/cnames-and-https-requirements.html>

<https://docs.aws.amazon.com/acm/latest/userguide/import-certificate.html>

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/CNAMEs.html>

upvoted 8 times

A developer is creating a template that uses AWS CloudFormation to deploy an application. The application is serverless and uses Amazon API Gateway, Amazon DynamoDB, and AWS Lambda.

Which AWS service or tool should the developer use to define serverless resources in YAML?

- A. CloudFormation serverless intrinsic functions
- B. AWS Elastic Beanstalk
- C. AWS Serverless Application Model (AWS SAM)
- D. AWS Cloud Development Kit (AWS CDK)

Correct Answer: C

Community vote distribution

C (100%)

 **Bibay**  1 year, 2 months ago

The recommended AWS service for defining serverless resources in YAML is the AWS Serverless Application Model (AWS SAM).

AWS SAM is an open-source framework that extends AWS CloudFormation to provide a simplified way to define the Amazon API Gateway APIs, AWS Lambda functions, and Amazon DynamoDB tables needed by your serverless application. You can define your serverless resources in a YAML template and then use the AWS SAM CLI to package and deploy your application.


AWS CloudFormation serverless intrinsic functions can also be used to define serverless resources in YAML, but they have some limitations compared to AWS SAM. AWS Elastic Beanstalk is a platform as a service (PaaS) that is not serverless specific, while the AWS Cloud Development Kit (AWS CDK) is an alternative to YAML-based templates that uses familiar programming languages like TypeScript, Python, and Java to define AWS infrastructure.

upvoted 20 times

 **jipark** 12 months ago

your explanation helps me a lot !

upvoted 2 times


 **Untamables**  1 year, 4 months ago

Selected Answer: C

C

<https://aws.amazon.com/serverless/sam/>

upvoted 6 times

 **sumanshu**  1 month, 1 week ago

Selected Answer: C

AWS SAM is specifically designed for serverless applications.

upvoted 1 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: C


==> Discard A. CloudFormation serverless intrinsic functions: Does not exist; CloudFormation has no intrinsic functions specifically for serverless.

==> Discard B. AWS Elastic Beanstalk: Not serverless; used for managing EC2-based applications.

==> Discard D. AWS Cloud Development Kit (AWS CDK): Does not use YAML; CDK defines infrastructure using programming languages.

C. AWS Serverless Application Model (AWS SAM): Extends CloudFormation, specifically designed for defining and deploying serverless resources (API Gateway, DynamoDB, Lambda) in YAML with simplified syntax. It perfectly aligns with the requirements in the question: serverless resources, YAML, and CloudFormation

upvoted 1 times

 **65703c1** 2 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗨️ 👤 **Jonalb** 9 months, 1 week ago

O AWS Serverless Application Model (AWS SAM) é uma extensão do AWS CloudFormation que facilita a definição de aplicações sem servidor. AWS SAM fornece modelos mais simples para configurar recursos sem servidor como AWS Lambda, Amazon API Gateway e Amazon DynamoDB. Os modelos podem ser definidos em YAML ou JSON.

C

upvoted 1 times

🗨️ 👤 **svrnvtr** 1 year, 4 months ago

Selected Answer: C

C is the answer

upvoted 3 times

A developer wants to insert a record into an Amazon DynamoDB table as soon as a new file is added to an Amazon S3 bucket. Which set of steps would be necessary to achieve this?

- A. Create an event with Amazon EventBridge that will monitor the S3 bucket and then insert the records into DynamoDB.
- B. Configure an S3 event to invoke an AWS Lambda function that inserts records into DynamoDB.
- C. Create an AWS Lambda function that will poll the S3 bucket and then insert the records into DynamoDB.
- D. Create a cron job that will run at a scheduled time and insert the records into DynamoDB.

Correct Answer: B

Community vote distribution

B (100%)

 **Bibay** Highly Voted 1 year, 2 months ago

The correct answer is B.

To insert a record into DynamoDB as soon as a new file is added to an S3 bucket, you can configure an S3 event notification to invoke an AWS Lambda function that inserts the records into DynamoDB. When a new file is added to the S3 bucket, the S3 event notification will trigger the Lambda function, which will insert the record into the DynamoDB table.

Option A is incorrect because Amazon EventBridge is not necessary to achieve this. S3 event notifications can directly invoke a Lambda function to insert records into DynamoDB.

Option C is incorrect because polling the S3 bucket periodically to check for new files is inefficient and not necessary with S3 event notifications.

Option D is incorrect because running a cron job at a scheduled time is not real-time and would not insert the record into DynamoDB as soon as a new file is added to the S3 bucket.

upvoted 13 times


 **Untamables** Highly Voted 1 year, 4 months ago

Selected Answer: B

B

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/NotificationHowTo.html>

upvoted 8 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B


A) Eliminated - This approach introduces unnecessary complexity when S3 already supports native event notifications to Lambda.

B) Correct - S3 has native support for event notifications to trigger Lambda functions

C) Eliminated - Polling is inefficient and unnecessary

D) Eliminated - This is a time-based solution, not event-driven, meaning there could be a delay between the file being added and the record being inserted.

upvoted 1 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: B

==> Discard A: EventBridge is unnecessary because S3 already provides direct event notifications.

==> Discard C: Polling violates the "as soon as" requirement due to delay and inefficiency.

==> Discard D: Cron jobs do not respond immediately, violating the "as soon as" requirement.

B is correct because S3 event notifications can trigger a Lambda function immediately when a new file is added, ensuring real-time insertion into DynamoDB.

B is better than A because S3 event notifications directly trigger a Lambda function, eliminating the need for additional configuration or services like EventBridge, making it simpler and more efficient.

upvoted 1 times

🗨️ 👤 **65703c1** 2 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗨️ 👤 **JohnPI** 6 months, 2 weeks ago

Selected Answer: B

A is also a solution for this which is better if we want loose coupling but will introduce a slight latency. The key word here is "as soon as" so the correct answer will be B.

upvoted 2 times

🗨️ 👤 **svrnvr** 1 year, 4 months ago

It is B

upvoted 4 times

A development team maintains a web application by using a single AWS CloudFormation template. The template defines web servers and an Amazon RDS database. The team uses the Cloud Formation template to deploy the Cloud Formation stack to different environments. During a recent application deployment, a developer caused the primary development database to be dropped and recreated. The result of this incident was a loss of data. The team needs to avoid accidental database deletion in the future. Which solutions will meet these requirements? (Choose two.)

- A. Add a CloudFormation Deletion Policy attribute with the Retain value to the database resource.
- B. Update the CloudFormation stack policy to prevent updates to the database.
- C. Modify the database to use a Multi-AZ deployment.
- D. Create a CloudFormation stack set for the web application and database deployments.
- E. Add a Cloud Formation DeletionPolicy attribute with the Retain value to the stack.

Correct Answer: AB

Community vote distribution

AB (100%)

 **Mtho96** Highly Voted 1 year ago

A. Add a CloudFormation Deletion Policy attribute with the Retain value to the database resource: By adding a DeletionPolicy attribute with the Retain value to the database resource in the CloudFormation template, the database will not be deleted even if the CloudFormation stack is deleted. This helps prevent accidental database loss during stack deletion.

B. Update the CloudFormation stack policy to prevent updates to the database: By updating the CloudFormation stack policy, the development team can restrict updates to the database resource. This prevents accidental modifications or recreations of the database during stack updates. The stack policy can define specific actions that are allowed or denied, providing an additional layer of protection against unintentional database changes.

upvoted 15 times


 **svrnvtr** Highly Voted 1 year, 4 months ago

Selected Answer: AB

AB

<https://aws.amazon.com/ru/premiumsupport/knowledge-center/cloudformation-accidental-updates/>

upvoted 7 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: AB

A) Correct - When a DeletionPolicy: Retain is applied to a resource, CloudFormation retains (does not delete) the resource even if the stack is deleted or updated.


B) Correct - Preventing updates ensures that no actions (such as deletion or replacement) can accidentally impact the database. This is an effective safeguard to avoid both deletion and unintended modifications.

C) Eliminated - Multi-AZ deployment is for ensuring database availability and durability, not to prevent accidental deletion. Even with Multi-AZ, the database can still be deleted.

D) Eliminated - Stack sets are for multi-account/multi-region deployments and do not inherently prevent resource deletion.

E) Eliminated - The DeletionPolicy attribute applies only to individual resources, not the entire stack.

upvoted 2 times

 **65703c1** 2 months, 1 week ago

Selected Answer: AB

AB is the correct answer.

upvoted 1 times

 **Vaibs099** 3 months, 1 week ago

A and B, A is straightforward. For B, within stack policy you can deny update to logical resource for RDS DB. This will prevent any updates to stack resource which could also erase and create new RDS instance.

upvoted 1 times

🗨️ 👤 **Jonalb** 9 months, 1 week ago

Selected Answer: AB

<https://aws.amazon.com/ru/premiumsupport/knowledge-center/cloudformation-accidental-updates/>

upvoted 1 times

🗨️ 👤 **magicjims** 10 months, 4 weeks ago

Selected Answer: AB

This came up in the exam today, I chose A&B

upvoted 3 times

🗨️ 👤 **panoptica** 10 months, 4 weeks ago

D & A for me

upvoted 2 times

🗨️ 👤 **nguyenta** 1 year ago

Selected Answer: AB

A and B

upvoted 2 times

🗨️ 👤 **marvel21** 1 year, 1 month ago

A & B Correct Answer

upvoted 2 times

🗨️ 👤 **s50600822** 1 year, 1 month ago

D because grandma said?

upvoted 2 times

🗨️ 👤 **Japanjot** 1 year, 2 months ago

A B CORRECT

upvoted 1 times

🗨️ 👤 **ihebchorfi** 1 year, 3 months ago

Selected Answer: AB

D is wrong, because while it still doesn't protect from the accidental deletion of the DB.

upvoted 1 times

🗨️ 👤 **ihebchorfi** 1 year, 3 months ago

After more thinking, combining A & D is the correct answer, so i would go with AD

upvoted 2 times

🗨️ 👤 **Untamables** 1 year, 4 months ago

Selected Answer: AB

A and B

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html>

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/protect-stack-resources.html>

upvoted 5 times

🗨️ 👤 **March2023** 1 year, 4 months ago

Selected Answer: AB

I agree it is AB

upvoted 3 times

A company has an Amazon S3 bucket that contains sensitive data. The data must be encrypted in transit and at rest. The company encrypts the data in the S3 bucket by using an AWS Key Management Service (AWS KMS) key. A developer needs to grant several other AWS accounts the permission to use the S3 GetObject operation to retrieve the data from the S3 bucket.

How can the developer enforce that all requests to retrieve the data provide encryption in transit?

- A. Define a resource-based policy on the S3 bucket to deny access when a request meets the condition "aws:SecureTransport": "false".
- B. Define a resource-based policy on the S3 bucket to allow access when a request meets the condition "aws:SecureTransport": "false".
- C. Define a role-based policy on the other accounts' roles to deny access when a request meets the condition of "aws:SecureTransport": "false".
- D. Define a resource-based policy on the KMS key to deny access when a request meets the condition of "aws:SecureTransport": "false".

Correct Answer: A

Community vote distribution

A (89%) 11%

 **Untamables** Highly Voted 1 year, 10 months ago


Selected Answer: A

A
<https://repost.aws/knowledge-center/s3-bucket-policy-for-config-rule>
 upvoted 9 times

 **Watascript** Highly Voted 1 year, 10 months ago

Selected Answer: A

A is correct.
 upvoted 5 times

 **sumanshu** Most Recent 1 month, 1 week ago


Selected Answer: A

A) Correct - The condition "aws:SecureTransport": "false" ensures that only secure requests (encrypted in transit) are allowed.


B) Eliminated - This allows access only when the request does not use secure transport. This is opposite of the requirement

C) Eliminated - Resource-based policies at the bucket level are better suited for cross-account access.

D) Eliminated - The GetObject operation is specific to the S3 bucket, not the KMS key.
 upvoted 2 times


 **rue_** 3 months ago

aws:SecureTransport condition does not apply to KMS key policies
 upvoted 1 times

 **65703c1** 8 months, 2 weeks ago

Selected Answer: A

A is the correct answer.
 upvoted 1 times

 **ibratoev** 10 months, 1 week ago

A is correct.
 upvoted 1 times

 **CrescentShared** 1 year, 3 months ago



Selected Answer: D

Hesitate between A and D.
 Question is not clear on weather we want to block all the information or only the sensitive part.
 upvoted 2 times

 **KarBiswa** 1 year, 1 month ago



Agree, but id we compare between A & D, A seems to be more accurate.

upvoted 1 times

  **winzzhhzzhh** 1 year, 5 months ago

I know A is correct but D seems correct as well, since users will need access to the KMS key to decrypt the data in the bucket.

upvoted 3 times

  **Malkia** 1 year, 9 months ago

Selected Answer: A

A is correct.

upvoted 1 times

An application that is hosted on an Amazon EC2 instance needs access to files that are stored in an Amazon S3 bucket. The application lists the objects that are stored in the S3 bucket and displays a table to the user. During testing, a developer discovers that the application does not show any objects in the list.

What is the MOST secure way to resolve this issue?

- A. Update the IAM instance profile that is attached to the EC2 instance to include the S3:* permission for the S3 bucket.
- B. Update the IAM instance profile that is attached to the EC2 instance to include the S3:ListBucket permission for the S3 bucket.
- C. Update the developer's user permissions to include the S3:ListBucket permission for the S3 bucket.
- D. Update the S3 bucket policy by including the S3:ListBucket permission and by setting the Principal element to specify the account number of the EC2 instance.

Correct Answer: B

Community vote distribution

B (80%)

A (20%)

🗳️ **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: B

The correct answer is B.

<https://repost.aws/knowledge-center/ec2-instance-access-s3-bucket>

Option A also works, but it is not compliant to the AWS security practice of the least privilege permissions.

upvoted 11 times

🗳️ **yeacuz** 1 year, 8 months ago

Option B only allows you to list the bucket - you will still not see the objects if only s3:ListBucket permission is configured.

upvoted 3 times

🗳️ **yeacuz** Highly Voted 1 year, 8 months ago

Selected Answer: A

Option A allows you to list buckets AND objects. Option B only allows you to list the bucket - you will still not see the objects if only s3:ListBucket permission is configured.

upvoted 5 times

🗳️ **nbxyzd** 3 months ago

Hey, don't mislead the others, please. Read the official document carefully before posting here.

upvoted 2 times

🗳️ **Jeremy11** 1 year, 6 months ago

Not true:

https://docs.aws.amazon.com/AmazonS3/latest/API/API_ListObjectsV2.html

To use this action in an AWS Identity and Access Management (IAM) policy, you must have permission to perform the s3:ListBucket action.

upvoted 2 times

🗳️ **ec8or** 10 months, 2 weeks ago

Answer is A: The question is not stating that the list of buckets cannot be seen but the objects within the lists cannot be seen. Seems the dev already has the s3:ListBucket option its the objects part that is missing.

upvoted 1 times

🗳️ **aaaaatoz** Most Recent 1 month ago

Selected Answer: B

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/access-policy-language-overview.html>

For example, the s3:ListBucket permission allows the user to use the Amazon S3 ListObjectsV2 operation. (The s3:ListBucket permission is a case where the action name doesn't map directly to the operation name.)

upvoted 1 times

🗳️ **sumanshu** 1 month, 1 week ago

Selected Answer: B

A) Eliminated - Too permissive: Grants more permissions than needed (S3:*), violating the principle of least privilege.

B) Correct - Minimal permissions: Only grants the permission needed for the application to perform the ListBucket operation. Permissions are tied to the EC2 instance's IAM instance profile, limiting access to that instance.

C) Eliminated - The developer's permissions are unrelated to the application running on the EC2 instance.

D) Eliminated - Grants permissions at the S3 bucket policy level, which applies to all resources in the account, not just the EC2 instance.

upvoted 1 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗨️ **Vaibs099** 9 months, 1 week ago

B is correct, Question is asking for lists the objects that are stored in the S3 bucket. s3:ListBucket gives bucket level objects list.

upvoted 1 times

🗨️ **ibratoev** 10 months, 1 week ago

The correct answer is B.

Option A works as well but only listing the files is mentioned as requirement.

upvoted 1 times

🗨️ **ninomfr64** 1 year, 5 months ago

Selected Answer: B

It is B, but I had to dig into docs to learn that to use ListObjectsV2, in an AWS Identity and Access Management (IAM) policy, you must have permission to perform the s3:ListBucket action.

https://docs.aws.amazon.com/AmazonS3/latest/API/API_ListObjectsV2.html

upvoted 2 times

🗨️ **jipark** 1 year, 6 months ago

are there anyone who can explain D ? - S3 bucket policy

upvoted 3 times

🗨️ **nmc12** 1 year, 4 months ago

Option D is not the most secure choice, as utilizing bucket policies and specifying account numbers can potentially lead to overly complex and less secure configurations, especially if not managed carefully.

To implement option B, follow these and it most secure!!!

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:ListBucket",
      "Resource": "arn:aws:s3:::your-bucket-name"
    }
  ]
}
```

upvoted 1 times

🗨️ **s50600822** 1 year, 7 months ago

A violated least privilege principle so B

upvoted 3 times

🗨️ **yashika2005** 1 year, 8 months ago

Selected Answer: B



the s3:ListBucket permission allows the user to use the Amazon S3 GET Bucket (List Objects) operation.

Reference: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/access-policy-language-overview.html>

upvoted 3 times

🗨️ **yashika2005** 1 year, 8 months ago

the s3:ListBucket permission allows the user to use the Amazon S3 GET Bucket (List Objects) operation.
Reference: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/access-policy-language-overview.html>
upvoted 1 times

  **svrnvtr** 1 year, 10 months ago

Selected Answer: B

It is B

upvoted 4 times

A company is planning to securely manage one-time fixed license keys in AWS. The company's development team needs to access the license keys in automaton scripts that run in Amazon EC2 instances and in AWS CloudFormation stacks.
Which solution will meet these requirements MOST cost-effectively?

- A. Amazon S3 with encrypted files prefixed with "config"
- B. AWS Secrets Manager secrets with a tag that is named SecretString
- C. AWS Systems Manager Parameter Store SecureString parameters
- D. CloudFormation NoEcho parameters

Correct Answer: C

Community vote distribution

C (100%)

🗳️ **alohayo** Highly Voted 10 months, 3 weeks ago

Both B and C are feasible solutions. Just consider the "MOST cost effectively" here.

AWS Systems Manager Parameter Store comes with no additional cost (Standard type). However, AWS Secrets Manager costs \$0.40 per secret per month, and data retrieval costs \$0.05 per 10,000 API calls.

C is much cheaper, guy.

upvoted 26 times

🗳️ **hanJR** Highly Voted 1 year, 3 months ago

I chose C because AWS Secrets Manager does auto key rotation(The question says that the key is one-time fixed).

upvoted 19 times

🗳️ **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: C

A) Eliminated - this solution involves additional complexity

B) Eliminated - Secrets Manager charges a fee for secret storage and API calls. For one-time fixed license keys, this might be unnecessary overhead. Secrets Manager is more suited for dynamically rotated secrets (e.g., database credentials), not fixed keys.

C) Correct: SecureString parameters are free to store unless using advanced features, making it highly cost-efficient.

D) Eliminated - Parameters marked with NoEcho are limited to the stack and cannot be accessed programmatically by EC2 instances

upvoted 1 times

🗳️ **Freddie26** 1 month, 1 week ago

Selected Answer: C

"One-time fixed license keys" are unique codes provided when purchasing a software license. It's more cost effective to provide this via AWS Systems Manager with secure parameters.

upvoted 1 times

🗳️ **65703c1** 2 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗳️ **william_cit** 4 months ago

It said 'one-time fixed license keys' and 'MOST cost-effectively', so C is better

upvoted 2 times

🗳️ **ibratoev** 4 months, 1 week ago

C seems the best fit.


upvoted 2 times

🗳️ **s50600822** 1 year, 1 month ago

PS prob is free for this use case <https://docs.aws.amazon.com/systems-manager/latest/userguide/parameter-store-advanced-parameters.html>, even though SM cost may also count to nothing(due to the scale of the use case and caching client).

Again the only notable difference is the aforementioned irrelevant tag.

upvoted 3 times

  **Untamables** 1 year, 4 months ago

Selected Answer: C

C

'<https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-parameter-store.html>

upvoted 9 times

A company has deployed infrastructure on AWS. A development team wants to create an AWS Lambda function that will retrieve data from an Amazon Aurora database. The Amazon Aurora database is in a private subnet in company's VPC. The VPC is named VPC1. The data is relational in nature. The Lambda function needs to access the data securely.

Which solution will meet these requirements?

- A. Create the Lambda function. Configure VPC1 access for the function. Attach a security group named SG1 to both the Lambda function and the database. Configure the security group inbound and outbound rules to allow TCP traffic on Port 3306.
- B. Create and launch a Lambda function in a new public subnet that is in a new VPC named VPC2. Create a peering connection between VPC1 and VPC2.
- C. Create the Lambda function. Configure VPC1 access for the function. Assign a security group named SG1 to the Lambda function. Assign a second security group named SG2 to the database. Add an inbound rule to SG1 to allow TCP traffic from Port 3306.
- D. Export the data from the Aurora database to Amazon S3. Create and launch a Lambda function in VPC1. Configure the Lambda function query the data from Amazon S3.



Correct Answer: A

Community vote distribution

A (64%)

C (29%)



7%

  **[Removed]**  1 year, 1 month ago

ooooh this one was rough. I am going with A --> <https://repost.aws/knowledge-center/connect-lambda-to-an-rds-instance>



I was between A and C... wording for both tricky. But the only way C would work is if the last portion of the sentence the read "Add an inbound rule to SG2 to allow TCP traffic from port 3306" or "Add an outbound rule to SG1 to allow TCP traffic..."

upvoted 9 times

  **9d8dd9c** 3 months, 2 weeks ago

But aren't the routing on SGs state-full so allowing inbound allows outbound too? or am I confusing that with something else?

upvoted 1 times

  **Yuri_024** 1 month, 2 weeks ago

SG1 is for the lambda, SG2 is for the database. In option C it says setting inbound traffic for the SG1 on TCP port 3306. But it should be setting inbound traffic for SG2 on TCP port 3306.

upvoted 2 times

  **shahs10**  1 year, 10 months ago

Selected Answer: A




Correct Answer is Answer A

For B creating new VPC for lambda does not seems a suitable solution

For C Assigning different security groups to both will not work

Option D will not be suitable for relational data and involve S3 in solution

upvoted 7 times

  **sumanshu**  1 month, 1 week ago



Selected Answer: A

B) Eliminated - Placing the Lambda function in a public subnet compromises security

C) Eliminated - The rule should allow traffic to SG2 (the database's security group) from SG1 (the Lambda function's security group), not the other way around.

D) Eliminated - Adds significant operational complexity

upvoted 1 times

  **wh1t4k3r** 5 months, 2 weeks ago

this one is badly written hehe

I would say A, but they missed to mention that this only works securely if the secgroup is listed as destination of the rules.

B would also work, but you need to properly configure it....

upvoted 1 times

🗨️ 👤 **Saurabh04** 6 months ago

Correction answer should be option C. Lambda function, configure VPC1 access, and assign separate security groups:

Lambda Function: Associate the Lambda function with VPC1.

Security Group (SG1): Assign SG1 to the Lambda function.

Security Group (SG2): Assign a second security group (SG2) to the Aurora database.

Inbound Rule: Add an inbound rule to SG1 to allow TCP traffic from Port 3306 (Aurora database port).

This approach ensures proper separation of concerns and simplifies security group management.

upvoted 1 times

🗨️ 👤 **tsangckl** 7 months, 2 weeks ago

This appear at 17 Jun exam

upvoted 4 times

🗨️ 👤 **65703c1** 8 months, 2 weeks ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ 👤 **ibratoev** 10 months, 1 week ago

A seems the answer, although a single SG for both the DB and Lambda is not a great practice. I would go with 2 SGs.

upvoted 3 times

🗨️ 👤 **TheFivePips** 11 months ago

Selected Answer: A

Security groups are statefull so you dont need to specify both inbound and outbound rules. However, you should have security groups on both resources as a best practice, and I dont think it is enough to have an inbound rule just on the lambda security group in this case.

This would essentially give the DB access to send traffic to the lambda function, rather than the lambda function accessing data from the DB like we want. If the lambda function doesnt have a permission on its security group letting it access the DB, then it will never communicate with it unless the DB contacts it first.

If C had placed the inbound permission on the DB, or if it had placed the outbound permission on the lambda then I think it would be right.

So while the wording is a little confusing, I think A is correct

upvoted 3 times

🗨️ 👤 **quanghao** 1 year, 3 months ago

Selected Answer: B

A Lambda function and RDS instance in different VPCs

First, use VPC peering to connect the two VPCs. Then, use the networking configurations to connect the Lambda function in one VPC to the RDS instance in the other:

upvoted 2 times

🗨️ 👤 **hcsaba1982** 1 year, 3 months ago

Selected Answer: B

This is the only one where lambda can reach the Database anyway, seems to me a prerequisite if the VPC was mentioned. Lambda by default, launched outside your VPC (in an AWS-owned VPC) so it cannot access resources.

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

if it were private maybe... but public so this answer definitely wrong

upvoted 1 times

🗨️ 👤 **dexdinh91** 1 year, 3 months ago

Selected Answer: B

B is correct?

upvoted 1 times

🗨️ 👤 **quanbui** 1 year, 3 months ago

Selected Answer: C

C, need 2 SG

upvoted 2 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

C the wording throws me off... Because the inbound rule in the end of the statement should be to the database not SG1. so we want to allow lambda access to the DB... The way this option is worded is not really giving lambda access to the db... it's giving DB access to lambda but not

the other way around which we need. So leaning with A

upvoted 1 times

🗨️ **sofiatian** 1 year, 4 months ago

Selected Answer: C

Need two security groups. One is for Lambda function. The other one is for DB

upvoted 1 times

🗨️ **konieczny69** 1 year ago

nonsense

why would anyone want sql application port access to lambda??

A is the only answer

upvoted 2 times

🗨️ **hsinchang** 1 year, 4 months ago

A. right

B. public, unsecure

C. excessive connections

D. additional cost and complexity

upvoted 3 times

🗨️ **love777** 1 year, 5 months ago

Selected Answer: A

VPC Configuration:

Ensure that your Lambda function is configured to run within the same VPC where your Amazon Aurora database resides (VPC1 in this case). Configure the Lambda function to use the appropriate subnets within VPC1, which are associated with the private subnet where your Amazon Aurora database is located.

Security Groups:

Attach a security group (SG1) to both the Lambda function and the Amazon Aurora database.

Configure the security group inbound rules for SG1 to allow incoming TCP traffic on Port 3306, which is the default port for MySQL (used by Aurora). This will allow communication between the Lambda function and the database.

Outbound rules should be allowed by default, so you don't need to make any changes there.

upvoted 2 times

🗨️ **ninomfr64** 1 year, 5 months ago

Selected Answer: A

There isn't the ideal solution to the use case among the options.

B) no need to create a new VPC and also you need to add route tables and configure SGs to make it works

C) this could work if the rule on SG1 was outbound instead of inbound (the connection is initiated from Lambda to Aurora)

D) export data to S3 is overkill and if you do that you no longer need to deploy the lambda in the VPC

A) works, as SG1 is attached to both Lambda and Aurora we need outbound rule to 3306 (Lambda initiate communication to Aurora) and also inbound rule from 3306 (to allow Aurora accept connection from Lambda). I don't like to have the same SG1 for both the Lambda and the Aurora

upvoted 5 times

A developer is building a web application that uses Amazon API Gateway to expose an AWS Lambda function to process requests from clients. During testing, the developer notices that the API Gateway times out even though the Lambda function finishes under the set time limit. Which of the following API Gateway metrics in Amazon CloudWatch can help the developer troubleshoot the issue? (Choose two.)

- A. CacheHitCount
- B. IntegrationLatency
- C. CacheMissCount
- D. Latency
- E. Count

Correct Answer: BD

Community vote distribution

BD (100%)

🗳️ **Untamables** Highly Voted 1 year, 4 months ago

Selected Answer: BD

B and D

The issue is caused by timeout. So the developer needs to know the latency information.

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-metrics-and-dimensions.html>

<https://repost.aws/knowledge-center/api-gateway-rest-api-504-errors>

upvoted 14 times

🗳️ **Watascript** Highly Voted 1 year, 4 months ago

Selected Answer: BD

<https://docs.aws.amazon.com/apigateway/latest/developerguide/monitoring-cloudwatch.html>

upvoted 5 times

🗳️ **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: BD

A) Eliminated - there is no mention of caching being involved in the problem, this metric is irrelevant for troubleshooting a timeout issue.

B) Correct - A high value for IntegrationLatency can indicate that the API Gateway is experiencing delays in receiving responses from Lambda

C) Eliminated - this metric is related to caching, which is not mentioned as part of the problem

D) Correct - high value for Latency can indicate where delays are occurring overall, including the Lambda function's processing time and any overhead in API Gateway.

E) Eliminated - While this metric provides information about the volume of requests, it does not help identify the cause of a timeout or latency issues.

upvoted 2 times

🗳️ **65703c1** 2 months, 1 week ago

Selected Answer: BD

BD is the correct answer.

upvoted 1 times

🗳️ **Jonalb** 9 months, 1 week ago

Selected Answer: BD

As melhores opções são, portanto, B. IntegraçãoLatência e D. Latência. Ambas as métricas fornecerão insights sobre onde pode estar ocorrendo a latência ou o atraso, ajudando o desenvolvedor a solucionar o problema.

upvoted 1 times

A development team wants to build a continuous integration/continuous delivery (CI/CD) pipeline. The team is using AWS CodePipeline to automate the code build and deployment. The team wants to store the program code to prepare for the CI/CD pipeline. Which AWS service should the team use to store the program code?

- A. AWS CodeDeploy
- B. AWS CodeArtifact
- C. AWS CodeCommit
- D. Amazon CodeGuru

Correct Answer: C

Community vote distribution

C (100%)

🗳️ **Untamables** Highly Voted 1 year, 4 months ago

Selected Answer: C

C
<https://aws.amazon.com/codecommit/>
upvoted 9 times

🗳️ **bahubba** Most Recent 4 weeks, 1 day ago

Selected Answer: C

This is likely outdated, as CodeCommit is being deprecated. They already do not allow new users to access CodeCommit.
upvoted 1 times

🗳️ **sumanshu** 1 month, 1 week ago

Selected Answer: C

CodeCommit is explicitly designed for storing and version-controlling program code,
upvoted 1 times

🗳️ **65703c1** 2 months, 1 week ago

Selected Answer: C

C is the correct answer.
upvoted 1 times

🗳️ **Lucian2407** 11 months, 1 week ago

Selected Answer: C

Simple answer: CodeCommit
upvoted 2 times

🗳️ **[Removed]** 7 months, 3 weeks ago

yep. I hope to get this one
upvoted 1 times

🗳️ **kgopireddy** 1 year, 4 months ago

Selected Answer: C

C is the right answer
upvoted 4 times

A developer is designing an AWS Lambda function that creates temporary files that are less than 10 MB during invocation. The temporary files will be accessed and modified multiple times during invocation. The developer has no need to save or retrieve these files in the future. Where should the temporary files be stored?

- A. the /tmp directory
- B. Amazon Elastic File System (Amazon EFS)
- C. Amazon Elastic Block Store (Amazon EBS)
- D. Amazon S3

Correct Answer: A

Community vote distribution

A (100%)

🗨️ **Untamables** Highly Voted 1 year, 4 months ago

Selected Answer: A

A

A Lambda function has access to local storage in the /tmp directory. Each execution environment provides between 512 MB and 10,240 MB, in 1-MB increments, of disk space in the /tmp directory.

<https://docs.aws.amazon.com/lambda/latest/dg/foundation-progmodel.html>

upvoted 18 times

🗨️ **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: A

A) Correct - AWS Lambda provides a /tmp directory with storage between 512B to 10,240 MB.

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-ephemeral-storage.html>. The files in this directory are ephemeral and will be deleted after the function execution ends.

B) Eliminated - Temporary files are not required to persist beyond the invocation, and using EFS would increase complexity and cost.

C) Eliminated - Lambda does not directly support EBS volumes

D) Eliminated- While S3 can store temporary data, it is not optimized for high-speed read/write operations during a Lambda invocation. Additionally, the use case specifies that the files do not need to be saved or retrieved later, making S3 an unnecessary overhead.

upvoted 2 times

🗨️ **65703c1** 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ **Mtho96** 1 year ago

The correct answer is A

The /tmp directory is the recommended location for storing temporary files within an AWS Lambda function. The /tmp directory provides a writable space with a local storage capacity of 512 MB. It is specifically designed for temporary storage within the Lambda execution environment.

upvoted 4 times

A developer is designing a serverless application with two AWS Lambda functions to process photos. One Lambda function stores objects in an Amazon S3 bucket and stores the associated metadata in an Amazon DynamoDB table. The other Lambda function fetches the objects from the S3 bucket by using the metadata from the DynamoDB table. Both Lambda functions use the same Python library to perform complex computations and are approaching the quota for the maximum size of zipped deployment packages.

What should the developer do to reduce the size of the Lambda deployment packages with the LEAST operational overhead?

- A. Package each Python library in its own .zip file archive. Deploy each Lambda function with its own copy of the library.
- B. Create a Lambda layer with the required Python library. Use the Lambda layer in both Lambda functions.
- C. Combine the two Lambda functions into one Lambda function. Deploy the Lambda function as a single .zip file archive.
- D. Download the Python library to an S3 bucket. Program the Lambda functions to reference the object URLs.

Correct Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **Untamables** Highly Voted 1 year, 4 months ago

Selected Answer: B

B

<https://docs.aws.amazon.com/lambda/latest/dg/invocation-layers.html>

upvoted 13 times

🗳️ 👤 **Ponyi** Highly Voted 9 months ago

Whenever you see "to make deployment package smaller" ----> Layers

upvoted 12 times

🗳️ 👤 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B

A) Eliminated - This approach would increase operational overhead because it requires packaging and maintaining duplicate libraries for each function.

B) Correct - A Lambda layer is a shared resource containing code or data that can be used by multiple Lambda functions. It reduces the size of individual deployment packages by offloading the library into a shared layer.

C) Eliminated - Combining the two functions increases code complexity

D) Eliminated - Each Lambda invocation would need to download the library, increasing cold-start time.

upvoted 1 times

🗳️ 👤 **65703c1** 2 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗳️ 👤 **Mtho96** 1 year ago

B

creating a Lambda layer with the required Python library and using it in both Lambda functions, is the most suitable solution for reducing the size of the deployment packages with minimal operational overhead.

<https://docs.aws.amazon.com/lambda/latest/dg/invocation-layers.html>

upvoted 6 times

🗳️ 👤 **Baba_Eni** 1 year, 1 month ago

Selected Answer: B

<https://docs.aws.amazon.com/lambda/latest/dg/invocation-layers.html>

upvoted 4 times

A developer is writing an AWS Lambda function. The developer wants to log key events that occur while the Lambda function runs. The developer wants to include a unique identifier to associate the events with a specific function invocation. The developer adds the following code to the Lambda function:

```
function handler(event, context) {
}
```

Which solution will meet this requirement?

- A. Obtain the request identifier from the AWS request ID field in the context object. Configure the application to write logs to standard output.
- B. Obtain the request identifier from the AWS request ID field in the event object. Configure the application to write logs to a file.
- C. Obtain the request identifier from the AWS request ID field in the event object. Configure the application to write logs to standard output.
- D. Obtain the request identifier from the AWS request ID field in the context object. Configure the application to write logs to a file.

Correct Answer: A

Community vote distribution

A (93%)

7%

 **ninomfr64** Highly Voted 11 months, 2 weeks ago


Selected Answer: A

Both A and D could work here, as both rely on the context object to get access to execution ID

https://docs.aws.amazon.com/us_en/lambda/latest/dg/python-context.html

While A uses stdout to send log to CloudWatch Log, D writes to a file. D is less specific (where is the file stored? A single file for each execution?) and looks more complex (manage file(s), manage concurrency access to the file ...), thus I'll go for A

upvoted 15 times

 **Untamables** Highly Voted 1 year, 4 months ago

Selected Answer: A

A

<https://docs.aws.amazon.com/lambda/latest/dg/nodejs-context.html>

<https://docs.aws.amazon.com/lambda/latest/dg/nodejs-logging.html>


There is no explicit information for the runtime, the code is written in Node.js.

upvoted 8 times

 **Pupina** 1 year ago

- <https://docs.aws.amazon.com/prescriptive-guidance/latest/implementing-logging-monitoring-cloudwatch/lambda-logging-metrics.html>
- Lambda automatically streams standard output and standard error messages from a Lambda function to CloudWatch Logs, without requiring logging drivers.

upvoted 2 times


 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: A

A) Correct - The context object contains the AWS request ID, which uniquely identifies each function invocation. Writing logs to standard output ensures they are automatically captured by CloudWatch Logs.

B) Eliminated - The event object contains input data passed to the Lambda function but does not include the AWS request ID. Writing logs to a file is not practical in Lambda because the filesystem is ephemeral, and logs stored in files would not persist beyond the function invocation.

upvoted 1 times

 **65703c1** 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ **badsati** 3 months, 3 weeks ago

Selected Answer: A

A should work
upvoted 1 times

🗨️ **james2033** 5 months, 1 week ago

Selected Answer: A

See `getAwsRequestId()` at <https://docs.aws.amazon.com/lambda/latest/dg/java-context.html>
upvoted 1 times

🗨️ **rimaSamir** 5 months, 2 weeks ago

Tricky question. Sure A and D both can do, but... The question is: why we need to get the request identifier if we will write logs to CloudWatch?
So, I will go with answer A.
upvoted 1 times

🗨️ **SD_CS** 6 months ago

I think it should be A. Also can anyone advise why the two answers are different ?
<https://www.examttopics.com/discussions/amazon/view/29007-exam-aws-certified-developer-associate-topic-1-question-26/>
upvoted 2 times

🗨️ **KarBiswa** 7 months, 2 weeks ago

The Option A is correct because:
The second argument is the context object. A context object is passed to your function by Lambda at runtime. This object provides methods and properties that provide information about the invocation, function, and runtime environment.
<https://docs.aws.amazon.com/lambda/latest/dg/python-handler.html>
upvoted 1 times

🗨️ **hsinchang** 10 months, 3 weeks ago

invocation is in the Context object, and logging into Standard output, which goes into CloudWatch(more durable, more scalable, etc.), is generally better than using temporary Files
upvoted 1 times

🗨️ **Pupina** 1 year ago

Selected Answer A:
Handler function <https://docs.aws.amazon.com/lambda/latest/dg/nodejs-handler.html>
Context object `awsRequestId` – The identifier of the invocation request. <https://docs.aws.amazon.com/lambda/latest/dg/nodejs-context.html>
upvoted 1 times

🗨️ **rInd2000** 1 year, 1 month ago

Selected Answer: A

In my opinion both options A and D can fulfill the requirement, since there is no requirement about any specific logging and monitoring tool I will go with defaults (A) because, simple is better than complex :)
upvoted 1 times

🗨️ **Prem28** 1 year, 2 months ago

Selected Answer: A

The application can write logs to standard output or to a file. Standard output is the default destination for logs. Logs that are written to standard output are sent to Amazon CloudWatch Logs. Logs that are written to a file are stored on the Lambda function's execution environment.
upvoted 5 times

🗨️ **Nagendhar** 1 year, 2 months ago

Ans: D

The code snippet provided in the question is obtaining the request identifier from the `context.awsRequestId` property, which is available in the context object provided to the Lambda function handler. Therefore, the correct option is:

D. Obtain the request identifier from the AWS request ID field in the context object. Configure the application to write logs to a file.

This option meets the requirement of logging key events and including a unique identifier to associate the events with a specific function invocation.

upvoted 1 times

🗨️ **Rpod** 1 year, 3 months ago

Selected Answer: D

Why not D ? Writing logs to a file seems more appropriate than stdout

upvoted 3 times



  **Watascript** 1 year, 4 months ago

Selected Answer: A

https://docs.aws.amazon.com/us_en/lambda/latest/dg/python-context.html

https://docs.aws.amazon.com/us_en/lambda/latest/dg/python-logging.html

upvoted 4 times

  **Dun6** 1 year, 4 months ago

Selected Answer: A

A it is

upvoted 3 times

A developer is working on a serverless application that needs to process any changes to an Amazon DynamoDB table with an AWS Lambda function.

How should the developer configure the Lambda function to detect changes to the DynamoDB table?

- A. Create an Amazon Kinesis data stream, and attach it to the DynamoDB table. Create a trigger to connect the data stream to the Lambda function.
- B. Create an Amazon EventBridge rule to invoke the Lambda function on a regular schedule. Conned to the DynamoDB table from the Lambda function to detect changes.
- C. Enable DynamoDB Streams on the table. Create a trigger to connect the DynamoDB stream to the Lambda function.
- D. Create an Amazon Kinesis Data Firehose delivery stream, and attach it to the DynamoDB table. Configure the delivery stream destination as the Lambda function.

Correct Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Untamables** Highly Voted 👍 1 year, 10 months ago

Selected Answer: C

C

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

upvoted 11 times

🗳️ 👤 **sumanshu** Most Recent 🕒 1 month, 1 week ago

Selected Answer: C

B) Eliminated - Polling the table for changes introduces unnecessary complexity and delay

C) Correct - DynamoDB Streams is specifically designed for this use case, capturing item-level changes in the table.

D) Eliminated - Firehose is used for data delivery and transformation to destinations like S3 or Redshift, not for processing DynamoDB table changes.

upvoted 1 times

🗳️ 👤 **65703c1** 8 months, 2 weeks ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗳️ 👤 **ibratoev** 10 months, 1 week ago

A should work as well. Why is it not A?

upvoted 1 times

🗳️ 👤 **nbxyzd** 3 months ago

While option A is technically possible, it introduces unnecessary complexity for detecting changes in a DynamoDB table. Using DynamoDB Streams directly is a more streamlined and optimized solution for this use case.

upvoted 3 times

🗳️ 👤 **nmc12** 1 year, 4 months ago

Selected Answer: C

C

Enabling DynamoDB Streams on the table allows you to capture and process changes (inserts, updates, deletes) to the table in real-time. You can then create a Lambda trigger that listens to the DynamoDB stream and invokes the Lambda function whenever there is a change in the table. This is a common and effective way to react to changes in DynamoDB tables with AWS Lambda functions.

upvoted 4 times

🗳️ 👤 **Baba_Eni** 1 year, 7 months ago

Selected Answer: C

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.html>
upvoted 2 times

An application uses an Amazon EC2 Auto Scaling group. A developer notices that EC2 instances are taking a long time to become available during scale-out events. The UserData script is taking a long time to run.

The developer must implement a solution to decrease the time that elapses before an EC2 instance becomes available. The solution must make the most recent version of the application available at all times and must apply all available security updates. The solution also must minimize the number of images that are created. The images must be validated.

Which combination of steps should the developer take to meet these requirements? (Choose two.)

- A. Use EC2 Image Builder to create an Amazon Machine Image (AMI). Install all the patches and agents that are needed to manage and run the application. Update the Auto Scaling group launch configuration to use the AMI.
- B. Use EC2 Image Builder to create an Amazon Machine Image (AMI). Install the latest version of the application and all the patches and agents that are needed to manage and run the application. Update the Auto Scaling group launch configuration to use the AMI.
- C. Set up AWS CodeDeploy to deploy the most recent version of the application at runtime.
- D. Set up AWS CodePipeline to deploy the most recent version of the application at runtime.
- E. Remove any commands that perform operating system patching from the UserData script.

Correct Answer: AC

Community vote distribution



imvb88 Highly Voted 1 year, 8 months ago

Selected Answer: AE

Why choose A over B? Problem is that B will tie an AMI with a specific version, so if there is a new version, we need to create a new AMI, and that contradicts with "minimize the number of images that are created".

Then E over C, D? E is obviously complementary to A, where removing commands from User Data will make the instance booting process much faster (and of course with A you don't need that anymore).

C and D also works but 1/not complementary with any other options; 2/CodeDeploy takes time to execute.

Hope this helps somebody struggling with this question.

upvoted 36 times

yashika2005 1 year, 8 months ago

thanksss a lott!

upvoted 1 times

minh12312312 1 year, 3 months ago

The solution must make the most recent version of the application available at all times

upvoted 4 times

[Removed] 1 year, 1 month ago

I agree I think between A and B.- answer is B

upvoted 5 times

r3mo 1 year, 6 months ago

And what about this requisit? "The solution must make the most recent version of the application available at all times". Only the Answer B fulfill this part.

upvoted 4 times

KillThemWithKindness Highly Voted 1 year, 6 months ago

Selected Answer: AC

Option E, which suggests removing operating system patching from the UserData script, might reduce the startup time. But this could leave your instances unpatched and vulnerable, which doesn't meet the requirement to apply all available security updates.

upvoted 15 times

maurice2005 10 months ago

well if u choose B and E then this will resolve as well

upvoted 2 times

🗨️ **Arad** Most Recent 1 month ago

Selected Answer: AE

AE is correct answer.

upvoted 1 times

🗨️ **sumanshu** 1 month, 1 week ago

Selected Answer: AC

B) Eliminated - Similar to option A, but also includes the latest version of the application in the AMI. While this reduces initialization time, it violates the requirement to "minimize the number of images" because it necessitates frequent AMI updates to keep the application version current.

C) CodeDeploy ensures that the most recent version of the application is deployed during runtime.

E) Eliminated - This violates the requirement to "apply all available security updates"

upvoted 1 times

🗨️ **Yuri_024** 1 month, 2 weeks ago

Selected Answer: AE

* The problem at hand is scaling-out events taking a long time (Instances need to spawn up fast, userData scripts must not run for long periods).

* CodePipeline or CodeDeploy does not play any role there.

* It does not mention an issue regarding the latest version of the application not being available.

* So I choose A & E

upvoted 1 times

🗨️ **youonebe** 1 month, 2 weeks ago

Selected Answer: AC

Option B is incorrect because:

Including the application in the AMI would require new image creation for every application update

Results in more frequent image builds than necessary

Increases maintenance overhead

upvoted 1 times

🗨️ **albert_kuo** 4 months, 2 weeks ago

Selected Answer: AC

B. While similar to A, this option includes installing the latest version of the application in the AMI. This is not ideal because it would require creating a new AMI every time the application is updated, which doesn't align with the requirement to minimize the number of images created.

upvoted 1 times

🗨️ **wh1t4k3r** 5 months, 2 weeks ago

Selected Answer: AC

B would require new images for each new app version. Idea is to minimize image creation, so A is a better fit.

E does speed up the process, BUT does not cover the app version requirement, nor the necessity to validate images, which codedeploy covers.

upvoted 2 times

🗨️ **MrDurian** 6 months, 2 weeks ago

IMO the correct answer is A and C.

Having a well set up AMI will reduce the need to run a long userData script.

Why not using B? Because that would couple the image with the app version. It is better to trigger a Code deploy that will deploy the latest version of the app on the 'optimized' AMI.

Regarding answer E, it would also be correct IMO but A and C seems to be the perfect matching scenario

upvoted 1 times

🗨️ **IYNH** 6 months, 2 weeks ago

Selected Answer: AC

The solution must make the most recent version of the application available at all times. B doesn't make sense because "latest version at the time AMI is created" becomes outdated when a newer one comes.

C is obviously needed to make the actual "latest" version deploy.

upvoted 2 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: BE

BE is the correct answer.

upvoted 1 times

🗨️ **MarcosSantos** 8 months, 3 weeks ago

I choose BE. Is better response

upvoted 1 times

🗨️ **ufuomaapoki** 9 months, 3 weeks ago

Selected Answer: AC

The requirements are:

1. Decreasing the time it takes for EC2 instances to become available during scale-out events.
2. Ensuring the most recent version of the application is available.
3. Applying all available security updates.
4. Minimising the number of images created.

[A] will satisfy requirements 1, 3, 4

[B] is similar to A, but will involve more AMI images

[C] Since the applications are on EC2 instances, CodeDeploy will do just fine to update the applications to the most recent version

[E] Removing any command for updates will leave our instances susceptible to vulnerabilities. Some commands can be removed, leaving the essential ones

upvoted 2 times

🗨️ **maurice2005** 10 months ago

B is faster than A.

E delegates all run time to AMI build time on B option.

upvoted 2 times

🗨️ **ibratoev** 10 months, 1 week ago

A and E.

A because number of images needs to be minimized.

E to speed up the boot time.

upvoted 1 times

🗨️ **KarBiswa** 10 months, 2 weeks ago

Selected Answer: BE

The most practical answers

upvoted 4 times

🗨️ **Abdullah22** 10 months, 3 weeks ago

Selected Answer: AC

going with ac

upvoted 1 times

A developer is creating an AWS Lambda function that needs credentials to connect to an Amazon RDS for MySQL database. An Amazon S3 bucket currently stores the credentials. The developer needs to improve the existing solution by implementing credential rotation and secure storage. The developer also needs to provide integration with the Lambda function.

Which solution should the developer use to store and retrieve the credentials with the LEAST management overhead?

- A. Store the credentials in AWS Systems Manager Parameter Store. Select the database that the parameter will access. Use the default AWS Key Management Service (AWS KMS) key to encrypt the parameter. Enable automatic rotation for the parameter. Use the parameter from Parameter Store on the Lambda function to connect to the database.
- B. Encrypt the credentials with the default AWS Key Management Service (AWS KMS) key. Store the credentials as environment variables for the Lambda function. Create a second Lambda function to generate new credentials and to rotate the credentials by updating the environment variables of the first Lambda function. Invoke the second Lambda function by using an Amazon EventBridge rule that runs on a schedule. Update the database to use the new credentials. On the first Lambda function, retrieve the credentials from the environment variables. Decrypt the credentials by using AWS KMS, Connect to the database.
- C. Store the credentials in AWS Secrets Manager. Set the secret type to Credentials for Amazon RDS database. Select the database that the secret will access. Use the default AWS Key Management Service (AWS KMS) key to encrypt the secret. Enable automatic rotation for the secret. Use the secret from Secrets Manager on the Lambda function to connect to the database.
- D. Encrypt the credentials by using AWS Key Management Service (AWS KMS). Store the credentials in an Amazon DynamoDB table. Create a second Lambda function to rotate the credentials. Invoke the second Lambda function by using an Amazon EventBridge rule that runs on a schedule. Update the DynamoDB table. Update the database to use the generated credentials. Retrieve the credentials from DynamoDB with the first Lambda function. Connect to the database.

Correct Answer: C

Community vote distribution

C (100%)

 **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: C


C

<https://docs.aws.amazon.com/secretsmanager/latest/userguide/intro.html>

https://docs.aws.amazon.com/secretsmanager/latest/userguide/create_database_secret.html

https://docs.aws.amazon.com/secretsmanager/latest/userguide/retrieving-secrets_lambda.html

upvoted 11 times

 **jjpark** 1 year, 6 months ago

"automatic rotation" "cross region" - Security Manager

upvoted 1 times

 **mooncake1** Most Recent 1 week ago

Selected Answer: C

A - KMS provides automatic rotation for AWS Managed Keys, but Custom Managed Keys are not provided (needs to implement rotation logic by dev). A did not specify which key the dev should use.

ANW : C - SM provides automatic rotation lambda for RDS Param rotation.

upvoted 1 times

 **sumanshu** 1 month, 1 week ago

Selected Answer: C

A) Eliminated - automatic credential rotation is limited

C) Correct - Secrets Manager is specifically designed for securely storing and retrieving secrets. Automatic rotation of credentials is fully supported



upvoted 1 times

 **Saudis** 3 months ago

Selected Answer: C

C The keyword automatic rotation for the parameter.



upvoted 1 times

  **65703c1** 8 months, 2 weeks ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

  **ibratoev** 10 months, 1 week ago

C. "credential rotation" = Security Manager



upvoted 2 times

  **jayvarma** 1 year, 5 months ago

Option C.

Keyword: Implementing credential rotation and secure storage.

upvoted 1 times

  **Mtho96** 1 year, 6 months ago

C

This solution minimizes management overhead by leveraging the built-in capabilities of AWS Secrets Manager, such as encryption, automatic rotation, and integration with AWS Lambda. It provides a secure and efficient way to store and retrieve

<https://docs.aws.amazon.com/secretsmanager/latest/userguide/intro.html>

https://docs.aws.amazon.com/secretsmanager/latest/userguide/create_database_secret.html

https://docs.aws.amazon.com/secretsmanager/latest/userguide/retrieving-secrets_Lambda.html

upvoted 2 times

A developer has written the following IAM policy to provide access to an Amazon S3 bucket:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:PutObject"
      ],
      "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
    },
    {
      "Effect": "Deny",
      "Action": "s3:*",
      "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET/secrets*"
    }
  ]
}
```


Which access does the policy allow regarding the s3:GetObject and s3:PutObject actions?

- A. Access on all buckets except the "DOC-EXAMPLE-BUCKET" bucket
- B. Access on all buckets that start with "DOC-EXAMPLE-BUCKET" except the "DOC-EXAMPLE-BUCKET/secrets" bucket
- C. Access on all objects in the "DOC-EXAMPLE-BUCKET" bucket along with access to all S3 actions for objects in the "DOC-EXAMPLE-BUCKET" bucket that start with "secrets"
- D. Access on all objects in the "DOC-EXAMPLE-BUCKET" bucket except on objects that start with "secrets"

Correct Answer: D

Community vote distribution

D (100%)


 **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: D

D

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/using-with-s3-actions.html>

upvoted 12 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: D


The condition restricts access to objects that start with the prefix secrets

upvoted 1 times

 **guidosolano** 7 months ago

Me toco en el examen


upvoted 1 times

 **65703c1** 8 months, 2 weeks ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

 **nmc12** 1 year, 4 months ago

Selected Answer: D

D

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/using-with-s3-actions.html>

upvoted 2 times

A developer is creating a mobile app that calls a backend service by using an Amazon API Gateway REST API. For integration testing during the development phase, the developer wants to simulate different backend responses without invoking the backend service. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an AWS Lambda function. Use API Gateway proxy integration to return constant HTTP responses.
- B. Create an Amazon EC2 instance that serves the backend REST API by using an AWS CloudFormation template.
- C. Customize the API Gateway stage to select a response type based on the request.
- D. Use a request mapping template to select the mock integration response.

Correct Answer: D

Community vote distribution

D (100%)

🗳️ **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: D

D

<https://docs.aws.amazon.com/apigateway/latest/developerguide/how-to-mock-integration.html>

upvoted 17 times

🗳️ **Dun6** Highly Voted 1 year, 10 months ago

Chatgpt said D

upvoted 6 times

🗳️ **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: D

C) Eliminated - Stages are not designed for selecting mock responses dynamically.

D) Correct - API Gateway provides a mock integration feature where request mapping templates can simulate responses.

upvoted 1 times

🗳️ **tsangckl** 7 months, 2 weeks ago

This appear at 17 Jun exam

upvoted 1 times

🗳️ **65703c1** 8 months, 2 weeks ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

🗳️ **Umuntu** 1 year, 3 months ago

D. Use a request mapping template to select the mock integration response.

Option D allows you to use a request mapping template in API Gateway to select the mock integration response. This approach allows you to simulate different backend responses without invoking the actual backend service. It provides flexibility and control over the responses without the need for additional AWS resources like Lambda functions or EC2 instances, thus minimizing operational overhead.

upvoted 5 times

🗳️ **hsinchang** 1 year, 4 months ago

without invoking backend service -> mock

upvoted 2 times

🗳️ **ninomfr64** 1 year, 5 months ago

Selected Answer: D

D as per doc <https://docs.aws.amazon.com/apigateway/latest/developerguide/how-to-mock-integration.html>

Wording confused me a bit, with mapping template you do not "select" a response, instead you actually craft it in this case

upvoted 1 times

🗨️ 👤 **KhyatiChhajer** 1 year, 8 months ago

Selected Answer: D

it's D

upvoted 1 times

🗨️ 👤 **March2023** 1 year, 10 months ago

Selected Answer: D

I'm going with D as well.

upvoted 4 times

A developer has a legacy application that is hosted on-premises. Other applications hosted on AWS depend on the on-premises application for proper functioning. In case of any application errors, the developer wants to be able to use Amazon CloudWatch to monitor and troubleshoot all applications from one place.

How can the developer accomplish this?

- A. Install an AWS SDK on the on-premises server to automatically send logs to CloudWatch.
- B. Download the CloudWatch agent to the on-premises server. Configure the agent to use IAM user credentials with permissions for CloudWatch.
- C. Upload log files from the on-premises server to Amazon S3 and have CloudWatch read the files.
- D. Upload log files from the on-premises server to an Amazon EC2 instance and have the instance forward the logs to CloudWatch.

Correct Answer: B

Community vote distribution

B (100%)

 **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: B

B

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/Install-CloudWatch-Agent.html>

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/install-CloudWatch-Agent-on-premise.html>


upvoted 14 times

 **Dun6** Highly Voted 1 year, 10 months ago

Selected Answer: B

We need cloudwatchagent

upvoted 6 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B

A) Eliminated - Using the AWS SDK means you must write custom code to capture and send logs to CloudWatch via the PutLogEvents API. This is not automatic as mentioned in option

B) CloudWatch agent is specifically designed for collecting logs and metrics from both AWS and on-premises environments.

C) Eliminated - unnecessary complexity.

D) Eliminated - unnecessarily complex and costly.


upvoted 3 times

 **Saudis** 3 months, 3 weeks ago

Selected Answer: B

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/Install-CloudWatch-Agent.html>


upvoted 2 times

 **65703c1** 8 months, 2 weeks ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

 **Baba_Eni** 1 year, 7 months ago

Selected Answer: B

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/Install-CloudWatch-Agent.html>

upvoted 2 times

An Amazon Kinesis Data Firehose delivery stream is receiving customer data that contains personally identifiable information. A developer needs to remove pattern-based customer identifiers from the data and store the modified data in an Amazon S3 bucket. What should the developer do to meet these requirements?

- A. Implement Kinesis Data Firehose data transformation as an AWS Lambda function. Configure the function to remove the customer identifiers. Set an Amazon S3 bucket as the destination of the delivery stream.
- B. Launch an Amazon EC2 instance. Set the EC2 instance as the destination of the delivery stream. Run an application on the EC2 instance to remove the customer identifiers. Store the transformed data in an Amazon S3 bucket.
- C. Create an Amazon OpenSearch Service instance. Set the OpenSearch Service instance as the destination of the delivery stream. Use search and replace to remove the customer identifiers. Export the data to an Amazon S3 bucket.
- D. Create an AWS Step Functions workflow to remove the customer identifiers. As the last step in the workflow, store the transformed data in an Amazon S3 bucket. Set the workflow as the destination of the delivery stream.

Correct Answer: A

Community vote distribution

A (100%)

🗃️ **Untamables** Highly Voted 1 year, 4 months ago

Selected Answer: A

A

<https://docs.aws.amazon.com/firehose/latest/dev/data-transformation.html>

upvoted 15 times

🗃️ **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: A

Kinesis Data Firehose supports the use of an AWS Lambda function for data transformation before delivering the data to a destination

upvoted 3 times

🗃️ **65703c1** 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗃️ **tttamttam** 1 year ago

Selected Answer: A

It supports custom data transformation using AWS Lambda

upvoted 3 times

A developer is using an AWS Lambda function to generate avatars for profile pictures that are uploaded to an Amazon S3 bucket. The Lambda function is automatically invoked for profile pictures that are saved under the /original/ S3 prefix. The developer notices that some pictures cause the Lambda function to time out. The developer wants to implement a fallback mechanism by using another Lambda function that resizes the profile picture.

Which solution will meet these requirements with the LEAST development effort?

- A. Set the image resize Lambda function as a destination of the avatar generator Lambda function for the events that fail processing.
- B. Create an Amazon Simple Queue Service (Amazon SQS) queue. Set the SQS queue as a destination with an on failure condition for the avatar generator Lambda function. Configure the image resize Lambda function to poll from the SQS queue.
- C. Create an AWS Step Functions state machine that invokes the avatar generator Lambda function and uses the image resize Lambda function as a fallback. Create an Amazon EventBridge rule that matches events from the S3 bucket to invoke the state machine.
- D. Create an Amazon Simple Notification Service (Amazon SNS) topic. Set the SNS topic as a destination with an on failure condition for the avatar generator Lambda function. Subscribe the image resize Lambda function to the SNS topic.

Correct Answer: A

Community vote distribution

A (65%)

C (19%)

B (16%)

🗨️ **March2023** Highly Voted 1 year, 10 months ago

Selected Answer: A

Wouldn't A be the Least Effort
upvoted 14 times

🗨️ **jingle4944** Highly Voted 1 year, 4 months ago

Selected Answer: A

Previously, you needed to write the SQS/SNS/EventBridge handling code within your Lambda function and manage retries and failures yourself.

With Destinations, you can route asynchronous function results as an execution record to a destination resource without writing additional code.

<https://aws.amazon.com/ru/blogs/compute/introducing-aws-lambda-destinations/>

upvoted 10 times

🗨️ **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: A

AWS Lambda supports Destinations for asynchronous invocations. You can configure a failure destination that triggers a fallback mechanism, such as invoking another Lambda function. This removes the need for additional services like SQS, SNS, or Step Functions, simplifying the architecture.

upvoted 2 times

🗨️ **f271c23** 1 month, 3 weeks ago

Selected Answer: B

Least development effort does not mean incomplete solution. what happens when the number of such issues are high and request needs to be queued. does not sound convincing

upvoted 1 times

🗨️ **lambdaFun** 3 months, 3 weeks ago

A is the correct Answer if we read the question in the last part "The developer wants to implement a fallback mechanism by using another Lambda function"

upvoted 2 times

🗨️ **PrinceMughal** 7 months, 1 week ago

Selected Answer: A

I will go with A as it is the simplest solution among other three option

upvoted 1 times

🗨️ **fabiomonta18** 7 months, 2 weeks ago

Selected Answer: B

This seems to be a tricky one. It's true that you can set Lambda destination, but you better set SQS as destination, exactly what the article suggests, go check it. The correct one is B.

<https://docs.aws.amazon.com/lambda/latest/dg/invoke-async.html#invoke-async-destinations>

upvoted 1 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ **vipyodha** 10 months, 1 week ago

B is right.

A.wrong because

Lambda Function Chaining: While Lambda function chaining is possible, it would require modifying the avatar generation Lambda function to include the resize function as a destination for failed events. This might involve additional coding and potentially more complex error handling within the Lambda function.

upvoted 1 times

🗨️ **Tluszczuk** 6 months, 2 weeks ago

With AWS Lambda destinations you don't need to interact with code to change this

upvoted 2 times

🗨️ **KarBiswa** 1 year, 1 month ago

Selected Answer: A

Least development effort no emphasis on orchestration

upvoted 3 times

🗨️ **KarBiswa** 11 months, 1 week ago

<https://aws.amazon.com/ru/blogs/compute/introducing-aws-lambda-destinations/> this link justifies the answer

upvoted 2 times

🗨️ **fabiomonta18** 7 months, 2 weeks ago

The article says "For each execution status such as Success or Failure you can choose one of four destinations: another Lambda function, SNS, SQS, or EventBridge." It hink for this reason the correct one is B.

upvoted 1 times

🗨️ **fabiomonta18** 7 months, 2 weeks ago

I mean, which destination other then SQS can you set in this case?

upvoted 1 times

🗨️ **Jonalb** 1 year, 3 months ago

Selected Answer: A

A. Defina a função Lambda de redimensionamento de imagem como um destino da função Lambda do gerador de avatar para os eventos que falham no processamento

upvoted 1 times

🗨️ **appuNBablu** 1 year, 4 months ago

A, because we can map another Lambda function as destination alongside (SQS, SNS, Event Bridge)

upvoted 1 times

🗨️ **ninomfr64** 1 year, 5 months ago

Selected Answer: A

A is the easiest option

<https://docs.aws.amazon.com/lambda/latest/dg/invoke-async.html#invoke-async-destinations>

upvoted 1 times

🗨️ **jayvarma** 1 year, 5 months ago

Option B is the right answer. Can someone say why B cannot be the right answer for this question?

Option A fails when there are huge amounts of requests coming to the lambda functions. There is every chance for lambda to throw ProvisionedThroughputExceeded Exception because of the throttling issues. Which is almost the similar reason why Option C will also fail at some point.

However, you could use SNS but it is not the best solution.

Definitely Option B.

upvoted 8 times

🗨️ 👤 **backfringe** 1 year, 6 months ago

Selected Answer: A

least amount of effort to set up destination on failure events to REsize Lambda

upvoted 1 times

🗨️ 👤 **AWSdeveloper08** 1 year, 6 months ago

Selected Answer: B

I agree with the explanation for option B. Scalability is the key

upvoted 2 times

🗨️ 👤 **[Removed]** 1 year, 6 months ago

Selected Answer: A

A is a simplest solution

<https://aws.amazon.com/ru/blogs/compute/introducing-aws-lambda-destinations/>

<https://docs.aws.amazon.com/lambda/latest/dg/invocation-async.html#invocation-async-destinations>

upvoted 3 times

🗨️ 👤 **jipark** 1 year, 6 months ago

your explanation looks correct.

Lambda "Denstination" seems exact solution for this.

it explains how to handle success, failed case.

upvoted 1 times

A developer needs to migrate an online retail application to AWS to handle an anticipated increase in traffic. The application currently runs on two servers: one server for the web application and another server for the database. The web server renders webpages and manages session state in memory. The database server hosts a MySQL database that contains order details. When traffic to the application is heavy, the memory usage for the web server approaches 100% and the application slows down considerably.

The developer has found that most of the memory increase and performance decrease is related to the load of managing additional user sessions. For the web server migration, the developer will use Amazon EC2 instances with an Auto Scaling group behind an Application Load Balancer.

Which additional set of changes should the developer make to the application to improve the application's performance?

- A. Use an EC2 instance to host the MySQL database. Store the session data and the application data in the MySQL database.
- B. Use Amazon ElastiCache for Memcached to store and manage the session data. Use an Amazon RDS for MySQL DB instance to store the application data.
- C. Use Amazon ElastiCache for Memcached to store and manage the session data and the application data.
- D. Use the EC2 instance store to manage the session data. Use an Amazon RDS for MySQL DB instance to store the application data.

Correct Answer: B

Community vote distribution

B (97%)

 **clarksu** Highly Voted 1 year, 4 months ago

Selected Answer: B

Option B ,

how can you image using an EC2 as cache

upvoted 10 times

 **Untamables** Highly Voted 1 year, 4 months ago

Selected Answer: B

B


Session stores are easy to create with Amazon ElastiCache for Memcached.

<https://aws.amazon.com/elasticache/memcached/>

With Amazon RDS, you can deploy scalable MySQL servers in minutes with cost-efficient and resizable hardware capacity.

<https://aws.amazon.com/rds/mysql/>

upvoted 10 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B

B) - in-memory data store for managing session state

C) Eliminated - ElastiCache is not a suitable replacement for a relational database like MySQL when dealing with structured data and complex queries.


upvoted 1 times

 **65703c1** 2 months, 1 week ago

Selected Answer: B

B is the correct answer.


upvoted 1 times

 **KarBiswa** 7 months, 2 weeks ago

Selected Answer: B

The additional requirement for the faster retrieval of data

upvoted 1 times

 **Aws_aspr** 11 months, 3 weeks ago

Selected Answer: B

B is correct

upvoted 1 times

🗨️ 👤 **nkelesidis** 1 year ago

Selected Answer: A

I choose A.

It says that the most of the memory increase is related to the load of managing additional user sessions. So I think Memcached doesn't make sense.

Also, isn't bad practice to store session information in db.

upvoted 1 times

🗨️ 👤 **ninomfr64** 11 months, 2 weeks ago

Session Store is one of the main use case for ElastiCache for Memcached as pwe AWS website

<https://aws.amazon.com/elasticache/memcached/#:~:text=ElastiCache%20for%20Memcached.,Session%20Store,-Session%20stores%20are>

upvoted 3 times

🗨️ 👤 **Dun6** 1 year, 4 months ago

Selected Answer: B

B it is

upvoted 6 times

An application uses Lambda functions to extract metadata from files uploaded to an S3 bucket; the metadata is stored in Amazon DynamoDB. The application starts behaving unexpectedly, and the developer wants to examine the logs of the Lambda function code for errors. Based on this system configuration, where would the developer find the logs?

- A. Amazon S3
- B. AWS CloudTrail
- C. Amazon CloudWatch
- D. Amazon DynamoDB

Correct Answer: C

Community vote distribution

C (100%)

🗉 **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: C

C

<https://docs.aws.amazon.com/prescriptive-guidance/latest/implementing-logging-monitoring-cloudwatch/lambda-logging-metrics.html>

upvoted 10 times

🗉 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: C

C) - AWS Lambda automatically sends logs generated by your function to Amazon CloudWatch Logs.

B) Eliminated - CloudTrail logs API activity and management operations across AWS, such as who invoked the Lambda function. However, it does not provide the function's execution logs.

upvoted 2 times

🗉 **tsangckl** 7 months, 2 weeks ago

This appear at 17 Jun exam

upvoted 4 times

🗉 **65703c1** 8 months, 2 weeks ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗉 **AhmedAliHashmi** 1 year, 5 months ago

Answer is C

upvoted 2 times

A company is using an AWS Lambda function to process records from an Amazon Kinesis data stream. The company recently observed slow processing of the records. A developer notices that the iterator age metric for the function is increasing and that the Lambda run duration is constantly above normal.

Which actions should the developer take to increase the processing speed? (Choose two.)

- A. Increase the number of shards of the Kinesis data stream.
- B. Decrease the timeout of the Lambda function.
- C. Increase the memory that is allocated to the Lambda function.
- D. Decrease the number of shards of the Kinesis data stream.
- E. Increase the timeout of the Lambda function.

Correct Answer: AC

Community vote distribution

AC (100%)

🗨️ **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: AC

A and C

<https://repost.aws/knowledge-center/lambda-iterator-age>

upvoted 15 times

🗨️ **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: AC

A) Kinesis uses shards to process data in parallel. If the data stream has a high volume of records, increasing the number of shards enables parallelism, allowing more Lambda instances to process data simultaneously

C) Allocating more memory to a Lambda function also increases the available CPU power

upvoted 2 times

🗨️ **nbxyzd** 3 months ago

AC

For those wondering why A is correct, quote:

Increasing the number of shards in a stream increases the number of concurrent Lambda functions consuming from your stream, which decreases a function's iterator age.

upvoted 2 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: AC

AC is the correct answer.

upvoted 1 times

🗨️ **KarBiswa** 1 year, 1 month ago

Selected Answer: AC

As the lambda has no timing issue

upvoted 1 times

🗨️ **gcmrjbr** 1 year, 2 months ago

CE

Shards (option A) works on the parallelism part and not on the function's execution time.

upvoted 1 times

🗨️ **gcmrjbr** 1 year, 2 months ago

A and C.

I would like to change my answer. More shards means more parallel processing.

upvoted 2 times

A company needs to harden its container images before the images are in a running state. The company's application uses Amazon Elastic Container Registry (Amazon ECR) as an image registry. Amazon Elastic Kubernetes Service (Amazon EKS) for compute, and an AWS CodePipeline pipeline that orchestrates a continuous integration and continuous delivery (CI/CD) workflow.


Dynamic application security testing occurs in the final stage of the pipeline after a new image is deployed to a development namespace in the EKS cluster. A developer needs to place an analysis stage before this deployment to analyze the container image earlier in the CI/CD pipeline.

Which solution will meet these requirements with the MOST operational efficiency?

- A. Build the container image and run the docker scan command locally. Mitigate any findings before pushing changes to the source code repository. Write a pre-commit hook that enforces the use of this workflow before commit.
- B. Create a new CodePipeline stage that occurs after the container image is built. Configure ECR basic image scanning to scan on image push. Use an AWS Lambda function as the action provider. Configure the Lambda function to check the scan results and to fail the pipeline if there are findings.
- C. Create a new CodePipeline stage that occurs after source code has been retrieved from its repository. Run a security scanner on the latest revision of the source code. Fail the pipeline if there are findings.
- D. Add an action to the deployment stage of the pipeline so that the action occurs before the deployment to the EKS cluster. Configure ECR basic image scanning to scan on image push. Use an AWS Lambda function as the action provider. Configure the Lambda function to check the scan results and to fail the pipeline if there are findings.

Correct Answer: B

Community vote distribution



Untamables Highly Voted 1 year, 4 months ago

Selected Answer: B

B

<https://docs.aws.amazon.com/AmazonECR/latest/userguide/image-scanning-basic.html>

The below blog post refers to the solution using Amazon Inspector and ECS, but the architecture is almost same as required in this scenario. The built in image scanning in Amazon ECR provides a simpler solution.

<https://aws.amazon.com/blogs/security/use-amazon-inspector-to-manage-your-build-and-deploy-pipelines-for-containerized-applications/>

upvoted 18 times

love777 Highly Voted 11 months, 1 week ago

Selected Answer: B

This approach integrates security scanning directly into the CI/CD pipeline and leverages AWS services for image scanning. Here's how it works:

A new CodePipeline stage is added after the container image is built, but before it's pushed to Amazon ECR.

ECR basic image scanning is configured to scan the image automatically upon push. This ensures that security scanning is part of the process.

An AWS Lambda function is used as an action provider in the pipeline. This Lambda function can be configured to analyze the scan results of the image.

If the Lambda function detects any security findings in the scan results, it can fail the pipeline, preventing the deployment of images with security vulnerabilities.

upvoted 10 times

sumanshu Most Recent 1 month, 1 week ago

Selected Answer: B

C) Eliminated - Scanning the source code directly does not analyze vulnerabilities specific to the container image

D) Eliminated - Adding the scan at the deployment stage is too late in the pipeline. It defeats the purpose of early detection and could allow vulnerable images to proceed further in the pipeline before being flagged

upvoted 2 times

🗨️ **trieudo** 1 month, 2 weeks ago

Selected Answer: B

B vs D:

B: scan before pushing ECR (after the container image is built)

D: scan after pushing ECR (before the deployment to the EKS cluster)

upvoted 1 times

🗨️ **65703c1** 2 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗨️ **ninomfr64** 11 months, 2 weeks ago

Selected Answer: B

B as per <https://docs.aws.amazon.com/amplify/latest/userguide/running-tests.html>

You can run end-to-end (E2E) tests in the test phase of your Amplify app to catch regressions before pushing code to production. The test phase can be configured in the build specification YAML. Currently, you can run only the Cypress testing framework during a build.

build specification is provided in the amplify.yml file

upvoted 2 times

🗨️ **imvb88** 1 year, 2 months ago

Selected Answer: D

So it narrows down to option B and D which using ECR basic image scanning.

B: create a stage

D: add an action to the existing stage

I'd go with D since executing an additional action will be faster than executing a whole stage.

upvoted 5 times

🗨️ **Toby_S** 1 year, 1 month ago

The question states "A developer needs to place an analysis stage" therefore I'd go with B.

upvoted 4 times

🗨️ **Rpod** 1 year, 3 months ago

Selected Answer: D

Chat GPT says D

upvoted 3 times

🗨️ **Umman** 1 year ago

ChatGPT says option B

upvoted 1 times

🗨️ **MrTee** 1 year, 3 months ago

Selected Answer: B

The developer should choose option B. Create a new CodePipeline stage that occurs after the container image is built. Configure ECR basic image scanning to scan on image push. Use an AWS Lambda function as the action provider. Configure the Lambda function to check the scan results and to fail the pipeline if there are findings. This will allow the developer to place an analysis stage before deployment to analyze the container image earlier in the CI/CD pipeline with the most operational efficiency.

CHATGPT

upvoted 5 times

A developer is testing a new file storage application that uses an Amazon CloudFront distribution to serve content from an Amazon S3 bucket. The distribution accesses the S3 bucket by using an origin access identity (OAI). The S3 bucket's permissions explicitly deny access to all other users.

The application prompts users to authenticate on a login page and then uses signed cookies to allow users to access their personal storage directories. The developer has configured the distribution to use its default cache behavior with restricted viewer access and has set the origin to point to the S3 bucket. However, when the developer tries to navigate to the login page, the developer receives a 403 Forbidden error.

The developer needs to implement a solution to allow unauthenticated access to the login page. The solution also must keep all private content secure.


Which solution will meet these requirements?

- A. Add a second cache behavior to the distribution with the same origin as the default cache behavior. Set the path pattern for the second cache behavior to the path of the login page, and make viewer access unrestricted. Keep the default cache behavior's settings unchanged.
- B. Add a second cache behavior to the distribution with the same origin as the default cache behavior. Set the path pattern for the second cache behavior to *, and make viewer access restricted. Change the default cache behavior's path pattern to the path of the login page, and make viewer access unrestricted.
- C. Add a second origin as a failover origin to the default cache behavior. Point the failover origin to the S3 bucket. Set the path pattern for the primary origin to *, and make viewer access restricted. Set the path pattern for the failover origin to the path of the login page, and make viewer access unrestricted.
- D. Add a bucket policy to the S3 bucket to allow read access. Set the resource on the policy to the Amazon Resource Name (ARN) of the login page object in the S3 bucket. Add a CloudFront function to the default cache behavior to redirect unauthorized requests to the login page's S3 URL.

Correct Answer: A

Community vote distribution

A (100%)

 **Untamables** Highly Voted 1 year, 4 months ago


Selected Answer: A

A

If you create additional cache behaviors, the default cache behavior is always the last to be processed.

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/distribution-web-values-specify.html#DownloadDistValuesCacheBehavior>

upvoted 12 times

 **sumanshu** Most Recent 1 month, 1 week ago


Selected Answer: A

A) Adding a second cache behavior allows you to define specific rules for the login page while keeping the default settings for private content unchanged.

B) Eliminated - Changing the default cache behavior to allow unrestricted access to the login page affects other private content, potentially compromising security.

D) Eliminated - Adding a bucket policy to allow public access to the login page directly contradicts the requirement to use CloudFront for secure content delivery.

upvoted 1 times

 **65703c1** 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

 **ShinobiGrappler** 7 months, 2 weeks ago

Answer is A. --The original way the developer had designed this application was too restrictive and didn't allow someone to even authenticate to get a signed cookie. By caching the second behavior, it allows the person authenticating to retrieve a cookie to access their personal data.

upvoted 1 times

🗨️ 👤 **LR2023** 8 months ago

D cloud front function acts as lamda function

upvoted 1 times

🗨️ 👤 **ninomfr64** 11 months, 2 weeks ago

Selected Answer: A

B) you cannot override the path pattern of the default Cache behavior

C) the origin failover is used when the primary origin is not available, this is not our case

D) with this configuration I think users wil get 403 Forbidden error and then redirected to the login page's S3 URL

A is a workable approach in my opinion

upvoted 2 times

🗨️ 👤 **Harddiver** 1 year, 1 month ago

Should it be D? In case s3 bucket restricts permissions, those should be open for login.

upvoted 3 times

🗨️ 👤 **MrTee** 1 year, 3 months ago

Selected Answer: A

By adding a second cache behavior with unrestricted viewer access to the login page's path pattern, unauthenticated users will be allowed to access the login page. At the same time, the default cache behavior's settings remain unchanged, and private content remains secure because it still requires signed cookies for access.

upvoted 3 times

A developer is using AWS Amplify Hosting to build and deploy an application. The developer is receiving an increased number of bug reports from users. The developer wants to add end-to-end testing to the application to eliminate as many bugs as possible before the bugs reach production.

Which solution should the developer implement to meet these requirements?

- A. Run the amplify add test command in the Amplify CLI.
- B. Create unit tests in the application. Deploy the unit tests by using the amplify push command in the Amplify CLI.
- C. Add a test phase to the amplify.yml build settings for the application.
- D. Add a test phase to the aws-exports.js file for the application.

Correct Answer: C

Community vote distribution

C (90%) 10%

🗳️ **gpt_test** Highly Voted 1 year, 10 months ago

Selected Answer: C

Explanation: Adding a test phase to the amplify.yml build settings allows the developer to define and execute end-to-end tests as part of the build and deployment process in AWS Amplify Hosting. This will help ensure that bugs are caught and fixed before the application reaches production, improving the overall quality of the application.

upvoted 16 times

🗳️ **Untamables** Highly Voted 1 year, 10 months ago

Selected Answer: C

C
<https://docs.aws.amazon.com/amplify/latest/userguide/running-tests.html>

upvoted 10 times

🗳️ **jipark** 1 year, 6 months ago

ton of thanks !!

document commented 'End to End Test'

upvoted 2 times

🗳️ **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: C

To implement end-to-end (E2E) testing before deployment, you can add a test phase to the amplify.yml file.

A) Eliminated - Amplify CLI does not have a command named amplify add test.

upvoted 1 times

🗳️ **tsangckl** 7 months, 2 weeks ago

This appear at 17 Jun exam

upvoted 1 times

🗳️ **65703c1** 8 months, 2 weeks ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗳️ **ninomfr64** 1 year, 5 months ago

Selected Answer: B

B as per <https://docs.aws.amazon.com/amplify/latest/userguide/running-tests.html>

You can run end-to-end (E2E) tests in the test phase of your Amplify app to catch regressions before pushing code to production. The test phase can be configured in the build specification YAML. Currently, you can run only the Cypress testing framework during a build.

build specification is provided in the amplify.yml file

upvoted 1 times

🗨️ 👤 **SachinR28** 1 year, 6 months ago

Selected Answer: B

I'LL GO WITH B

upvoted 1 times

🗨️ 👤 **rind2000** 1 year, 8 months ago

Selected Answer: B

We can use amplify.yml file to run any test commands at build time. Since the test must run while the program is being deployed (E2E) I'll go with B.

upvoted 1 times

An ecommerce company is using an AWS Lambda function behind Amazon API Gateway as its application tier. To process orders during checkout, the application calls a POST API from the frontend. The POST API invokes the Lambda function asynchronously. In rare situations, the application has not processed orders. The Lambda application logs show no errors or failures.

What should a developer do to solve this problem?

- A. Inspect the frontend logs for API failures. Call the POST API manually by using the requests from the log file.
- B. Create and inspect the Lambda dead-letter queue. Troubleshoot the failed functions. Reprocess the events.
- C. Inspect the Lambda logs in Amazon CloudWatch for possible errors. Fix the errors.
- D. Make sure that caching is disabled for the POST API in API Gateway.


Correct Answer: B

Community vote distribution

B (53%)

A (31%)

Other

  **[Removed]**  8 months, 2 weeks ago

B. Create and inspect the Lambda dead-letter queue. Troubleshoot the failed functions. Reprocess the events.

When Lambda functions are invoked asynchronously, there is a possibility that the function fails without logging errors if there is no proper error handling. Configuring a dead-letter queue (DLQ) allows you to capture and review events that were not processed successfully. By creating and inspecting the DLQ, you can identify and troubleshoot the issues with the failed Lambda invocations and reprocess those events if needed. This will help ensure that no orders are missed without leaving any errors in the Lambda logs.

upvoted 23 times

  **gpt_test**  1 year, 10 months ago

Selected Answer: B

Explanation: By configuring a dead-letter queue (DLQ) for the Lambda function, you can capture asynchronous invocation events that were not successfully processed. This allows you to troubleshoot the failed functions and reprocess the events, ensuring that orders are not missed. The DLQ will hold information about the failed events, allowing you to analyze and resolve the issue.

upvoted 13 times

  **rind2000** 1 year, 8 months ago

as you said "... events that were not successfully processed." but there is not failure in Lambda log, so the lambda was not invoked by the POST API event. B is id not the answer.

upvoted 5 times

  **kavi00203** 1 year, 7 months ago

Its an asynchronous invocation events, that's y there is no log.




Because in asynchronous its not mandatory to get the result after invocation events.

upvoted 2 times

  **TeeTheMan** 1 year, 6 months ago

Asynchronous invocation means that the caller of the lambda does not wait for a response. The type of invocation has no effect on the lambda having logs or not. I picked A, because the lambda not having logs suggests something's gone wrong upstream of the lambda.

upvoted 5 times

  **sumanshu**  1 month, 1 week ago

Selected Answer: B

since there are no logs of errors or failures, the issue is likely due to dropped asynchronous events. Configuring a DLQ will capture these dropped events for further analysis and reprocessing.

A) Eliminated - While inspecting frontend logs can help diagnose client-side issues, the problem here involves unprocessed events in Lambda.

C) Eliminated - The question states that there are no errors or failures in the Lambda logs

D) Eliminated - API Gateway caching is not related to this issue. Even with caching enabled, events would still invoke the Lambda function.

upvoted 2 times

🗨️ **mallikarjun_angadi** 2 months, 2 weeks ago

Answer : B

Lambda has a concurrency limit and in some cases, if the limit is reached, Lambda could throttle incoming requests without throwing an error, which means some invocations may be lost or delayed. Those will be moved to DLQ

upvoted 1 times

🗨️ **nbxyzd** 3 months ago

Selected Answer: B

Option A is obviously wrong. Remember, it's an asynchronous lambda, so replaying the POST API returns no info instrumental to issue diagnosis. If you know how AWS usually designs a quiz, you'll know B is definitely the answer.

upvoted 1 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗨️ **KarBiswa** 1 year, 1 month ago

Selected Answer: B

<https://aws.amazon.com/about-aws/whats-new/2016/12/aws-lambda-supports-dead-letter-queues/>

upvoted 1 times

🗨️ **Jonalb** 1 year, 3 months ago

Selected Answer: B

B. Crie e inspecione a fila de mensagens mortas do Lambda. Solucione os problemas das funções com falha. Reprocesse os eventos. Mais Votados

upvoted 1 times

🗨️ **mr_swal** 1 year, 3 months ago

Selected Answer: A

The Lambda application logs show no errors or failures. - So Lambda function was not invoked at all

upvoted 1 times

🗨️ **daicoso** 1 year, 3 months ago

if the application code doesn't log errors and doesn't throw exceptions, no error or failure will be logged

upvoted 1 times

🗨️ **nmc12** 1 year, 4 months ago

Selected Answer: B

The Lambda Dead Letter Queue is a feature that helps in troubleshooting events that failed processing by a Lambda function. When an asynchronous invocation of a Lambda function fails, AWS Lambda can direct the failed event to an Amazon SNS topic or an Amazon SQS queue (the dead-letter queue), where the event is stored and can be analyzed or reprocessed.

upvoted 1 times

🗨️ **norris81** 1 year, 4 months ago

Selected Answer: C

I don't like B which has reprocess the errors, which will make a whole load of errors be process creating orders which could be months old

upvoted 3 times

🗨️ **misa27** 1 year, 4 months ago

Selected Answer: B

B

<https://aws.amazon.com/what-is/dead-letter-queue/>

upvoted 1 times

🗨️ **ninomfr64** 1 year, 5 months ago

Selected Answer: B

A) asynchronous invocations do not return result to the caller, thus I do not expect errors in frontend log

C) the scenario question rules out the option to have error messages in the Lambda log

D) I do not see how caching can have impact in this scenario

B) having a dead-letter queue is a viable option to troubleshoot asynchronous lambda invocation error, another option would be using Destination

upvoted 1 times

🗨️ **backfringe** 1 year, 6 months ago

Selected Answer: C

Option C is the appropriate choice because it involves inspecting the Lambda logs in Amazon CloudWatch to identify any potential issues or errors that might be causing the orders not to be processed

Option B is not the most appropriate choice because the dead-letter queue is generally used to capture events that cannot be processed by a Lambda function. In this scenario, it seems that the Lambda function is executing without apparent errors. Thus, the issue might not be related to dead-letter queue failures.

upvoted 2 times

🗨️ **redfivedog** 1 year, 6 months ago

Selected Answer: D

I think D should be the correct answer to this question. The logs have no indications of errors or failed events, so if some transactions are not being processed, that probably means that the lambda function wasn't invoked for those calls. One reason could be that caching is enabled in API gateway for the POST request, so the lambda function isn't triggered for any cache hits.

- A is not correct as the frontend would be getting 202s for all asynchronous post requests.

- B is not correct because lambda logs have no errors => no lambda execution errors => DLQ won't get any requests of interest if we enable it. A comment below mentioned that asynchronous lambda invocations don't generate logs, but that is not true.

- C is obviously incorrect. The premise explicitly mentions that there aren't any errors in the logs.

upvoted 3 times

🗨️ **Saurabh04** 5 months, 2 weeks ago

Caching is only for GET Requests not for POST Requests. Correct answer is B

upvoted 2 times

🗨️ **xdkonorek2** 1 year, 1 month ago

Absolutely agree, D is the answer

upvoted 1 times

🗨️ **gomurali** 1 year, 7 months ago

<https://aws.amazon.com/about-aws/whats-new/2016/12/aws-lambda-supports-dead-letter-queues/>

upvoted 1 times

🗨️ **csG13** 1 year, 8 months ago

Selected Answer: B

It's B. Apparently C & D are wrong.

Also it's not A because the call is async. Meaning that the response code from the lambda service is 202. Since generally frontend can make POST requests, the problem should be visible somewhere in the backed. Dead-letter queues are for debugging and further analysis. Hence should be B.

upvoted 3 times

🗨️ **m5357** 1 year, 4 months ago

How can you tell from this context that the POST API call was successful?

upvoted 1 times

A company is building a web application on AWS. When a customer sends a request, the application will generate reports and then make the reports available to the customer within one hour. Reports should be accessible to the customer for 8 hours. Some reports are larger than 1 MB. Each report is unique to the customer. The application should delete all reports that are older than 2 days. Which solution will meet these requirements with the LEAST operational overhead?

- A. Generate the reports and then store the reports as Amazon DynamoDB items that have a specified TTL. Generate a URL that retrieves the reports from DynamoDB. Provide the URL to customers through the web application.
- B. Generate the reports and then store the reports in an Amazon S3 bucket that uses server-side encryption. Attach the reports to an Amazon Simple Notification Service (Amazon SNS) message. Subscribe the customer to email notifications from Amazon SNS.
- C. Generate the reports and then store the reports in an Amazon S3 bucket that uses server-side encryption. Generate a presigned URL that contains an expiration date. Provide the URL to customers through the web application. Add S3 Lifecycle configuration rules to the S3 bucket to delete old reports.
- D. Generate the reports and then store the reports in an Amazon RDS database with a date stamp. Generate a URL that retrieves the reports from the RDS database. Provide the URL to customers through the web application. Schedule an hourly AWS Lambda function to delete database records that have expired date stamps.

Correct Answer: C

Community vote distribution

C (100%)

gpt_test Highly Voted 1 year, 3 months ago

Selected Answer: C

Explanation: Storing the reports in an Amazon S3 bucket provides a cost-effective and scalable solution for handling files larger than 1 MB. Server-side encryption ensures data security. Generating a presigned URL with an expiration date allows the customer to access the report for 8 hours, and S3 Lifecycle configuration rules automatically delete the reports older than 2 days, reducing operational overhead.
upvoted 15 times

March2023 Highly Voted 1 year, 4 months ago

Selected Answer: C

Presigned URL
upvoted 8 times

sumanshu Most Recent 1 month, 1 week ago

Selected Answer: C

S3 Lifecycle rules can automatically delete objects after 2 days, reducing operational overhead by eliminating the need for custom cleanup logic

A) Eliminated - DynamoDB is not designed to store large objects (e.g., reports larger than 1 MB).

D) Eliminated - Storing large reports in a relational database like RDS is inefficient and costly compared to S3.
upvoted 1 times

sumanshu 1 month, 1 week ago

B) Eliminated - Attaching reports to SNS messages is impractical because SNS is not designed for large file attachments.
upvoted 1 times

65703c1 2 months, 1 week ago

Selected Answer: C

C is the correct answer.
upvoted 1 times

KarBiswa 7 months, 2 weeks ago

Selected Answer: C

The 1MB condition denies the TTL option so C is best
upvoted 2 times

LR2023 8 months ago

C

presigned and lifecycle rules to move
upvoted 1 times

🗨️ 👤 **ninomfr64** 11 months, 2 weeks ago

A) DynamoDB cannot store object larger than 400K

B) SNS cannot send email with attachment - <https://repost.aws/questions/QUOvaKJVb3QzOqVENONBZUag/sns-send-file-attachment>

D) the nature or format of the report is not specified, however RDS doesnt look like a great place to store large document file. Also generating a url to the reports from the RDS database requires some work while it is a native capabilities in S3

C) is a workable solution as S3 is designed to store file objects, it allows to easily generate pre-signed url, and provide lifecycle management rule that allows to expire objects

upvoted 6 times

🗨️ 👤 **imvb88** 1 year, 2 months ago

Selected Answer: C

Dynamo DB cannot store object > 400KB -> option A is out immediately.

Limited access to S3 calls for presigned URL which is option C. C also has lifecycle config to delete old object while B does not have that.

D is possible but too much effort compared to design pattern in C.

upvoted 5 times

🗨️ 👤 **Untamables** 1 year, 4 months ago

Selected Answer: C

C

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/using-presigned-url.html>

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lifecycle-mgmt.html>

upvoted 5 times


A company has deployed an application on AWS Elastic Beanstalk. The company has configured the Auto Scaling group that is associated with the Elastic Beanstalk environment to have five Amazon EC2 instances. If the capacity is fewer than four EC2 instances during the deployment, application performance degrades. The company is using the all-at-once deployment policy.

What is the MOST cost-effective way to solve the deployment issue?

- A. Change the Auto Scaling group to six desired instances.
- B. Change the deployment policy to traffic splitting. Specify an evaluation time of 1 hour.
- C. Change the deployment policy to rolling with additional batch. Specify a batch size of 1.
- D. Change the deployment policy to rolling. Specify a batch size of 2.

Correct Answer: C

Community vote distribution



gpt_test Highly Voted 1 year, 3 months ago

Selected Answer: C

Explanation: The rolling with additional batch deployment policy allows Elastic Beanstalk to launch additional instances in a new batch before terminating the old instances. In this case, specifying a batch size of 1 means that Elastic Beanstalk will deploy the application updates to 1 new instance at a time, ensuring that there are always at least 4 instances available during the deployment process. This method maintains application performance while minimizing the additional cost.

upvoted 19 times

gagol14 Highly Voted 1 year, 1 month ago

Selected Answer: C

1. Rolling with additional batch deployment: This type of deployment maintains full capacity while new application versions are deployed. It launches a new batch of instances with the new application version, and if the new batch is healthy, it terminates a batch of instances running the old application version.

2. Batch size of 1: This will ensure that one new instance is launched with the new version of the application. Once it is deemed healthy, one of the old instances will be terminated. This will continue until all instances are running the new version, ensuring the capacity is never less than four instances. This approach will add only a minimal additional cost for the temporary overlapping instances during deployment.

upvoted 13 times

sumanshu Most Recent 1 month, 1 week ago

Selected Answer: C

A) Eliminated - Permanently increasing the desired capacity to six EC2 instances incurs unnecessary ongoing costs.

B) Eliminated - Traffic splitting is primarily used for canary testing or gradual traffic shifting. It does not guarantee maintaining the desired instance capacity during deployment.

D) Eliminated - A rolling deployment with a batch size of 2 means two instances are replaced at a time. With a fixed capacity of 5 instances, this would temporarily reduce the number of running instances to 3, causing performance degradation.

upvoted 2 times

sumanshu 1 month, 1 week ago

C) Batch Size of 1: Specifies that only one instance will be updated at a time, minimizing downtime and maintaining the required instance capacity for performance.

Rolling with additional Batch - This deployment strategy adds an extra batch of EC2 instances during deployment. After updating one batch, the additional instances are terminated, ensuring cost efficiency.

upvoted 1 times

65703c1 2 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗨️ 👤 **Aleam** 7 months, 1 week ago

Selected Answer: D

Option D is the best solution because it allows the company to update the application without losing service or reducing availability significantly, and without increasing the cost or complexity of the solution.

upvoted 2 times

🗨️ 👤 **poklakni** 3 months, 2 weeks ago

the requirement is not to go below 4 instances. Option D specifies a batch size of 2 which would lead to 3 running instances. The correct option is C

upvoted 1 times

🗨️ 👤 **KarBiswa** 7 months, 2 weeks ago

Selected Answer: C

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.rolling-version-deploy.html>

upvoted 2 times

🗨️ 👤 **quangphungdev218** 1 year ago

Selected Answer: C

The correct answer is: C

upvoted 1 times

🗨️ 👤 **Prem28** 1 year, 1 month ago

The correct answer is: D. Change the deployment policy to rolling. Specify a batch size of 2.

A rolling deployment policy will deploy the new application version to one batch of instances at a time, while the other batches continue to serve traffic. This ensures that the application always has at least four instances available during the deployment.

Specifying a batch size of 2 means that two instances will be deployed at a time. This is the most cost-effective option because it minimizes the number of instances that are needed to maintain application performance during the deployment.

The other options are not as cost-effective because they require more instances to be running during the deployment. Option A requires six instances, option B requires at least five instances, and option C requires at least four instances.

upvoted 2 times

🗨️ 👤 **gagol14** 1 year, 1 month ago

The rolling deployment policy updates a few instances at a time, but unlike the "rolling with additional batch" option, it does not launch new instances before terminating the old ones. Therefore, capacity could drop below four during deployment, affecting application performance.

upvoted 2 times

🗨️ 👤 **jipark** 1 year ago

C: cost 1 additional EC2

D : degrade performance

it looks exam gave key "2 batch" meaning - do not choose this answer.

upvoted 1 times

🗨️ 👤 **nmc12** 10 months ago

If batch size of 1:

During the time the new instances are being deployed and are not yet in service, there are only $5 - 2 = 3$ old instances available to serve the traffic, which violates the requirement to maintain at least 4 instances to avoid performance degradation.

so, i go with A answer.

upvoted 1 times

🗨️ 👤 **Untamables** 1 year, 4 months ago

Selected Answer: C

C

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.rolling-version-deploy.html>

upvoted 4 times

A developer is incorporating AWS X-Ray into an application that handles personal identifiable information (PII). The application is hosted on Amazon EC2 instances. The application trace messages include encrypted PII and go to Amazon CloudWatch. The developer needs to ensure that no PII goes outside of the EC2 instances.

Which solution will meet these requirements?

- A. Manually instrument the X-Ray SDK in the application code.
- B. Use the X-Ray auto-instrumentation agent.
- C. Use Amazon Macie to detect and hide PII. Call the X-Ray API from AWS Lambda.
- D. Use AWS Distro for Open Telemetry.

Correct Answer: A

Community vote distribution

A (82%)


B (18%)

 **gpt_test** Highly Voted 1 year, 3 months ago

Selected Answer: A

Explanation: By manually instrumenting the X-Ray SDK in the application code, the developer can have full control over which data is included in the trace messages. This way, the developer can ensure that no PII is sent to X-Ray by carefully handling the PII within the application and not including it in the trace messages.

upvoted 25 times

 **Untamables** Highly Voted 1 year, 4 months ago

Selected Answer: A

A

Not to send any PII to AWS X-Ray service, add instrumentation code in your application at each location to send trace information that PII is eliminated.

<https://docs.aws.amazon.com/xray/latest/devguide/xray-instrumenting-your-app.html>

upvoted 7 times

 **Niluka** Most Recent 1 month ago

Selected Answer: D

Currently, with the all-at-once deployment policy, Elastic Beanstalk updates all EC2 instances simultaneously. This can cause issues with application performance, as the entire fleet of instances is replaced at once, leading to potential downtime or degraded performance if the number of available instances falls below the necessary threshold (fewer than four EC2 instances in your case).

Switching to a rolling deployment with a batch size of 2 will allow Elastic Beanstalk to update the instances in smaller batches. During each batch update, only two instances will be updated at a time, which means that there will always be a sufficient number of EC2 instances running (at least three instances at all times). This avoids the performance degradation caused by having fewer than four EC2 instances available.

upvoted 1 times

 **sumanshu** 1 month, 1 week ago

Selected Answer: A

C) Eliminated - Amazon Macie is designed for identifying and securing sensitive data stored in AWS services like S3

A) A (manual instrumentation) is the best solution because it gives the developer full control over what data is sent to X-Ray, ensuring that no PII leaves the EC2 instances.

B) Eliminated - Auto-instrumentation automatically captures data without offering granular control over what is sent to X-Ray. This approach could inadvertently send sensitive PII data to X-Ray, violating the requirements.

upvoted 1 times

 **65703c1** 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ 👤 **TheFivePips** 4 months, 4 weeks ago

Selected Answer: A

X-Ray auto-instrumentation agent itself does not inherently remove or redact Personally Identifiable Information (PII). The primary purpose of the auto-instrumentation agent is to automate the process of instrumenting supported frameworks and libraries for tracing with AWS X-Ray.

When dealing with PII or any sensitive information, the responsibility for ensuring that such data is not exposed in traces lies with the application code and configuration, rather than the X-Ray auto-instrumentation agent.

While the X-Ray auto-instrumentation agent simplifies the instrumentation process, the need for precise control over PII and the ability to implement custom security measures make manual instrumentation more suitable in this scenario.

upvoted 2 times

🗨️ 👤 **SerialIDr** 5 months, 1 week ago

Selected Answer: A

A. To ensure that no personally identifiable information (PII) goes outside of the EC2 instances while incorporating AWS X-Ray into an application that handles PII, the developer should manually instrument the X-Ray SDK in the application code. This approach allows for precise control over what data is captured and sent to X-Ray, enabling the developer to exclude or anonymize PII before it leaves the application environment, thereby meeting the requirement to ensure that no PII goes outside of the EC2 instances.

upvoted 1 times

🗨️ 👤 **SerialIDr** 6 months, 3 weeks ago

Selected Answer: A

This approach allows for granular control over what data is captured and sent to AWS X-Ray. The developer can instrument the code to ensure that PII is either not included in the trace data or is properly encrypted before being sent. This method provides the necessary control to meet the requirement.

upvoted 1 times

🗨️ 👤 **a_win** 7 months, 1 week ago

Selected Answer: B

The X-Ray auto-instrumentation agent can help ensure that sensitive information like PII is not transmitted outside of the EC2 instances. It automatically instruments the application without requiring manual intervention, making it easier to maintain traceability without risking the exposure of sensitive data.

Options A and D involve manual or custom instrumentations, which might inadvertently expose PII if not implemented correctly. Option C, using Amazon Macie to detect and hide PII and calling the X-Ray API from Lambda, might add complexity to the architecture and doesn't directly address the prevention of PII leaving the EC2 instances.

upvoted 2 times

🗨️ 👤 **chewasa** 7 months, 3 weeks ago

Selected Answer: B

Option B, using the X-Ray auto-instrumentation agent, is the most appropriate solution for ensuring that no PII goes outside of the EC2 instances.

upvoted 5 times

🗨️ 👤 **chewasa** 7 months, 3 weeks ago

A. Manually instrumenting the X-Ray SDK in the application code might lead to the possibility of inadvertently including PII in trace messages, and it may not be as foolproof as the auto-instrumentation agent.

B. The X-Ray auto-instrumentation agent automatically instruments the supported runtime environments, making it less error-prone and ensuring that sensitive information like PII is not leaked.

upvoted 1 times

🗨️ 👤 **chewasa** 7 months, 3 weeks ago

C. Amazon Macie is a service designed for discovering, classifying, and protecting sensitive data, but using it to detect and hide PII in combination with X-Ray is not a standard approach. It's more focused on data discovery and classification.

D. AWS Distro for OpenTelemetry is an observability project but may not provide the same level of automation for ensuring that no PII goes outside of the EC2 instances as the X-Ray auto-instrumentation agent.

upvoted 1 times

🗨️ 👤 **love777** 11 months, 1 week ago

Selected Answer: B

The X-Ray auto-instrumentation agent is designed to automatically trace and collect data from AWS resources and services without requiring manual instrumentation in your application code.

It helps ensure that sensitive information, such as PII, remains within the EC2 instances by not transmitting the data outside explicitly. The agent focuses on tracing the application behavior and performance without directly sending PII to external services.

This solution is suitable for ensuring compliance and data security while still benefiting from X-Ray's tracing and insights.

upvoted 2 times

🗨️ 👤 **r3mo** 1 year ago

Option "B" : Because. Avoids human error.

upvoted 2 times

🗨️ 👤 **Umman** 1 year ago

Using the X-Ray auto-instrumentation agent (Option B) is the best choice in this scenario because it will automatically instrument the application without requiring any manual code changes. Additionally, when using X-Ray with auto-instrumentation, you can control the sampling rate to ensure that only a subset of trace data (and encrypted PII) is sent to X-Ray and CloudWatch, reducing the risk of sensitive data being exposed outside of the instances.

upvoted 2 times

🗨️ 👤 **jasper_pigeon** 1 year ago

For non-Java applications running on EC2 instances, you will need to use the appropriate X-Ray SDKs to manually instrument the application code. You can't use auto-agent

upvoted 2 times

🗨️ 👤 **kris_je**c 1 year ago

Its very clear from Macie definition that it also provides automated protection as well apart from findings the PII data

upvoted 1 times

🗨️ 👤 **ttamttam** 1 year ago

Selected Answer: A

I think B is incorrect as the auto instrument cannot hide it, right?

upvoted 1 times

🗨️ 👤 **dan80** 1 year, 3 months ago

Selected Answer: A

C is wrong, Amazon Macie discover PII but dont hide it


upvoted 3 times

A developer is migrating some features from a legacy monolithic application to use AWS Lambda functions instead. The application currently stores data in an Amazon Aurora DB cluster that runs in private subnets in a VPC. The AWS account has one VPC deployed. The Lambda functions and the DB cluster are deployed in the same AWS Region in the same AWS account. The developer needs to ensure that the Lambda functions can securely access the DB cluster without crossing the public internet. Which solution will meet these requirements?

- A. Configure the DB cluster's public access setting to Yes.
- B. Configure an Amazon RDS database proxy for the Lambda functions.
- C. Configure a NAT gateway and a security group for the Lambda functions.
- D. Configure the VPC, subnets, and a security group for the Lambda functions.

Correct Answer: D

Community vote distribution



jayvarma Highly Voted 1 year, 5 months ago

Option D is the right answer. When we want the lambda to privately access the DB cluster instead of moving the traffic over the public internet, we need to have the lambda and db cluster to be in the same VPC.

When we configure the VPC, subnets, and a security group for the lambda function, the lambda function will be able to communicate with the db cluster using the private IPs that are associated to the VPC.

NAT gateway comes into use when you have the lambda deployed in a private subnet and you would want to provide internet access to it.
upvoted 17 times

gpt_test Highly Voted 1 year, 10 months ago

Selected Answer: D

Explanation: To securely access the Amazon Aurora DB cluster without crossing the public internet, the Lambda functions need to be configured to run within the same VPC as the DB cluster. This involves configuring the VPC, subnets, and a security group for the Lambda functions. This setup ensures that the Lambda functions can communicate with the DB cluster using private IP addresses within the VPC.

upvoted 8 times

sumanshu Most Recent 1 month, 1 week ago

Selected Answer: D

Lambda functions can be configured to run within a VPC. By assigning the Lambda functions to the same VPC and private subnets as the Aurora DB cluster, the communication remains internal to the VPC and does not cross the public internet. Configuring the security group ensures that the Lambda functions can securely connect to the Aurora DB cluster by allowing appropriate inbound/outbound rules.

upvoted 1 times

AmitRanchi 3 months, 3 weeks ago

Selected Answer: D

Ans is D.

upvoted 1 times

65703c1 8 months, 2 weeks ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

Wendy1113 1 year, 2 months ago

B

<https://repost.aws/questions/QUXXSqEPGbQx6qiyBa1D1Udg/lambda-to-db-connectivity-best-practices>

upvoted 2 times

maurice2005 10 months, 1 week ago

Actually Proxy should be on the same VPC as the database and since lambda is in another vpc it doesnt have access unless a connection happens between these two vpc or just option D

upvoted 1 times

🗨️ **alex_heavy** 1 year, 4 months ago

Selected Answer: B

<https://www.udemy.com/course/aws-certified-developer-associate-dva-c01/learn/lecture/36527788#overview>

<https://aws.amazon.com/ru/blogs/compute/using-amazon-rds-proxy-with-aws-lambda/>

upvoted 1 times

🗨️ **eberhe900** 1 year, 7 months ago

Selected Answer: C

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-vpc.html>

upvoted 2 times

🗨️ **ejlp** 6 months, 1 week ago

After reading doc in the link that you mention, my conclusion is D

NAT GW is required if Lamba in the user VPC need to access internet

upvoted 1 times

🗨️ **Untamables** 1 year, 10 months ago

Selected Answer: D

D

<https://docs.aws.amazon.com/lambda/latest/dg/foundation-networking.html>

upvoted 5 times

🗨️ **Dun6** 1 year, 10 months ago

Selected Answer: D

D is correct, NATGateway is for when we want Lambda to access the public when it is in a private VPC

upvoted 6 times

A developer is building a new application on AWS. The application uses an AWS Lambda function that retrieves information from an Amazon DynamoDB table. The developer hard coded the DynamoDB table name into the Lambda function code. The table name might change over time. The developer does not want to modify the Lambda code if the table name changes.

Which solution will meet these requirements MOST efficiently?

- A. Create a Lambda environment variable to store the table name. Use the standard method for the programming language to retrieve the variable.
- B. Store the table name in a file. Store the file in the /tmp folder. Use the SDK for the programming language to retrieve the table name.
- C. Create a file to store the table name. Zip the file and upload the file to the Lambda layer. Use the SDK for the programming language to retrieve the table name.
- D. Create a global variable that is outside the handler in the Lambda function to store the table name.

Correct Answer: A

Community vote distribution



A (100%)

  **Dun6** Highly Voted 1 year, 4 months ago

Selected Answer: A

You need to use environment variables

upvoted 9 times

  **Untamables** Highly Voted 1 year, 4 months ago

Selected Answer: A

A

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-envvars.html>



upvoted 7 times

  **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: A

Environment variables are a built-in feature of AWS Lambda



upvoted 1 times

  **65703c1** 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

  **mma34** 8 months, 2 weeks ago

Selected Answer: A

Why are some answers wrong on here?

upvoted 1 times

  **eberhe900** 1 year ago

Selected Answer: A

You can use environment variables to adjust your function's behavior without updating code. An environment variable is a pair of strings that is stored in a function's version-specific configuration. The Lambda runtime makes environment variables available to your code and sets additional environment variables that contain information about the function and invocation request.

upvoted 2 times

  **gpt_test** 1 year, 3 months ago

Selected Answer: A

Explanation: Using Lambda environment variables allows you to store configuration information separate from your code, which makes it easy to update the table name without changing the Lambda function code. AWS Lambda provides built-in support for environment variables, making it the most efficient solution.

upvoted 5 times

A company has a critical application on AWS. The application exposes an HTTP API by using Amazon API Gateway. The API is integrated with an AWS Lambda function. The application stores data in an Amazon RDS for MySQL DB instance with 2 virtual CPUs (vCPUs) and 64 GB of RAM.

Customers have reported that some of the API calls return HTTP 500 Internal Server Error responses. Amazon CloudWatch Logs shows errors for "too many connections." The errors occur during peak usage times that are unpredictable.


The company needs to make the application resilient. The database cannot be down outside of scheduled maintenance hours.

Which solution will meet these requirements?

- A. Decrease the number of vCPUs for the DB instance. Increase the max_connections setting.
- B. Use Amazon RDS Proxy to create a proxy that connects to the DB instance. Update the Lambda function to connect to the proxy.
- C. Add a CloudWatch alarm that changes the DB instance class when the number of connections increases to more than 1,000.
- D. Add an Amazon EventBridge rule that increases the max_connections setting of the DB instance when CPU utilization is above 75%.

Correct Answer: B

Community vote distribution



MrTee **Highly Voted** 1 year, 3 months ago

Selected Answer: B

The best solution to meet these requirements would be to use Amazon RDS Proxy to create a proxy that connects to the DB instance and update the Lambda function to connect to the proxy.

upvoted 11 times

sumanshu **Most Recent** 1 month, 1 week ago

Selected Answer: B

A) Eliminated - Decreasing the number of vCPUs will worsen performance.

B) When your Lambda function wants to query the database, instead of opening a new connection every time (which is what causes "too many connections"), it talks to the RDS Proxy. For example, if 100 Lambda functions try to connect at the same time, instead of opening 100 database connections, the RDS Proxy might only open 10 and reuse them.

C) Eliminated - It incurs unnecessary costs for increased capacity, even if the issue is caused by connection management.

upvoted 1 times

65703c1 2 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

SerialiDr 6 months, 3 weeks ago

Selected Answer: B

Amazon RDS Proxy is designed to handle a large number of simultaneous connections efficiently. It sits between your application and your RDS database to pool and share database connections, improving database efficiency and application scalability. This approach can reduce the number of connections to the database and handle unpredictable peak loads more effectively.

upvoted 3 times

hsinchang 10 months, 3 weeks ago

Selected Answer: B

B: RDS Proxy establishes and manages the necessary connection pools to your database so that your Lambda function creates fewer database connections¹. RDS Proxy also handles failovers and retries automatically, which improves the availability of your application.

A will reduce the performance and capacity of the database.

C may incur additional charges for scaling up the DB instance. It may also cause downtime during the scaling process, which violates the requirement that the database cannot be down outside of scheduled maintenance hours.

D may not react fast enough to handle unpredictable peak usage times. It may also cause memory issues if the `max_connections` setting is too high.

upvoted 1 times

🗨️ 👤 **love777** 11 months ago

Selected Answer: B

Adding an Amazon EventBridge rule to increase the `max_connections` setting based on CPU utilization is not directly addressing the issue of too many connections. Additionally, focusing solely on CPU utilization might not be the best metric for handling connection-related issues.

upvoted 2 times

🗨️ 👤 **tttamttam** 1 year ago

Selected Answer: B

I think D is incorrect because it increases the number of connections based on the CPU consumption not the number of connections.

upvoted 1 times

🗨️ 👤 **Naj_64** 1 year ago

Selected Answer: D

<https://repost.aws/knowledge-center/rds-mysql-max-connections>

upvoted 1 times

🗨️ 👤 **csG13** 1 year, 1 month ago

Selected Answer: B

It's B. RDS proxy can handle many open connections to the database.

upvoted 3 times

🗨️ 👤 **awsdummie** 1 year, 2 months ago

Selected Answer: D

There should not be any downtime. Create an Event bridge rule to update the `max_connections` parameter in Parameter group of DB instance.

upvoted 1 times

A company has installed smart meters in all its customer locations. The smart meters measure power usage at 1-minute intervals and send the usage readings to a remote endpoint for collection. The company needs to create an endpoint that will receive the smart meter readings and store the readings in a database. The company wants to store the location ID and timestamp information.

The company wants to give its customers low-latency access to their current usage and historical usage on demand. The company expects demand to increase significantly. The solution must not impact performance or include downtime while scaling.

Which solution will meet these requirements MOST cost-effectively?

- A. Store the smart meter readings in an Amazon RDS database. Create an index on the location ID and timestamp columns. Use the columns to filter on the customers' data.
- B. Store the smart meter readings in an Amazon DynamoDB table. Create a composite key by using the location ID and timestamp columns. Use the columns to filter on the customers' data.
- C. Store the smart meter readings in Amazon ElastiCache for Redis. Create a SortedSet key by using the location ID and timestamp columns. Use the columns to filter on the customers' data.
- D. Store the smart meter readings in Amazon S3. Partition the data by using the location ID and timestamp columns. Use Amazon Athena to filter on the customers' data.

Correct Answer: B

Community vote distribution

B (100%)

MrTee **Highly Voted** 1 year, 3 months ago

Selected Answer: B

The most cost-effective solution to meet these requirements would be to store the smart meter readings in an Amazon DynamoDB table and create a composite key using the location ID and timestamp columns
upvoted 9 times

sumanshu **Most Recent** 1 month, 1 week ago

Selected Answer: B

DynamoDB is purpose-built for low-latency, scalable storage of high-frequency, time-series data.
Composite key design (location ID + timestamp) enables efficient querying.
Automatically scales without downtime or performance impact.
upvoted 1 times

65703c1 2 months, 1 week ago

Selected Answer: B

B is the correct answer.
upvoted 1 times

SerialiDr 6 months, 3 weeks ago

Selected Answer: B

This solution provides low-latency access to real-time and historical data, scales seamlessly to accommodate increased demand without downtime, and is likely to be more cost-effective than the alternatives for this specific use case. DynamoDB's managed service nature also reduces the administrative burden of managing the database.
upvoted 2 times

Gold07 9 months, 2 weeks ago

C is the right answer
upvoted 2 times


Naj_64 1 year ago

Selected Answer: B

Going with B. DynamoDB is the most cost-effective solution.
upvoted 3 times

  **jasper_pigeon** 1 year ago

You need to use Athena as well to do partitoning
upvoted 2 times

  **HuiHsin** 1 year, 1 month ago

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-sort-keys.html>
upvoted 1 times

A company is building a serverless application that uses AWS Lambda functions. The company needs to create a set of test events to test Lambda functions in a development environment. The test events will be created once and then will be used by all the developers in an IAM developer group. The test events must be editable by any of the IAM users in the IAM developer group.

Which solution will meet these requirements?

- A. Create and store the test events in Amazon S3 as JSON objects. Allow S3 bucket access to all IAM users.
- B. Create the test events. Configure the event sharing settings to make the test events shareable.
- C. Create and store the test events in Amazon DynamoDB. Allow access to DynamoDB by using IAM roles.
- D. Create the test events. Configure the event sharing settings to make the test events private.

Correct Answer: B

Community vote distribution

B (80%)

A (18%)

 **renekton** Highly Voted 1 year, 9 months ago

Selected Answer: B

Under the "Test" tab there's an option. (Shareable)

This event is available to IAM users within the same account who have permissions to access and use shareable events.

You can check this by yourself on the Lambda

Also, here's a documentation

<https://docs.aws.amazon.com/lambda/latest/dg/testing-functions.html#creating-shareable-events>


upvoted 30 times

 **delak** Highly Voted 1 year, 8 months ago

Selected Answer: B

Since March of this year, this is now possible to share test events within the same account with different users.


upvoted 6 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B

(Shareable events) directly leverages AWS Lambda's built-in functionality, reducing the need for custom storage and retrieval mechanisms.


upvoted 2 times

 **Anandesh** 6 months, 2 weeks ago

Selected Answer: B

<https://docs.aws.amazon.com/lambda/latest/dg/testing-functions.html>

upvoted 1 times

 **65703c1** 8 months, 2 weeks ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

 **SerialiDr** 11 months, 1 week ago

Selected Answer: A

This option is the most straightforward and aligns with AWS practices for managing shared resources like test events. IAM policies can be configured to grant the necessary permissions to the developer group, ensuring that all members can access and edit the test events stored in S3. This method leverages the scalability and security features of S3, along with the granular permission control provided by IAM, to meet the requirements.

upvoted 1 times

 **manngupta007** 1 year ago

Answer: B

<https://aws.amazon.com/about-aws/whats-new/2022/03/aws-lambda-console-test-events/>

upvoted 1 times

SerialiDr 1 year ago

Selected Answer: A

This option is viable. Amazon S3 can store JSON objects (test events), and access to these objects can be controlled through S3 bucket policies or IAM policies. By setting the correct permissions, all IAM users in the developer group can read and write to the S3 bucket, enabling them to edit and use the test events.

upvoted 1 times

ez_24 1 year ago

Selected Answer: A

The key Concept here is Sharing - test events in the Lambda console are for individual account can't be used by other developers

upvoted 1 times

a_win 1 year, 1 month ago

Selected Answer: A

This approach ensures that the test events are stored centrally in an S3 bucket where all IAM users within the developer group have access. By granting access to the S3 bucket to all IAM users, any user within the group can create, edit, and retrieve the test events, meeting the requirement for collaborative access and editing.

Options B and D don't directly address the need for IAM users to edit the test events; sharing settings might allow access, but they might not allow editing by all IAM users in the group. Option C, using DynamoDB, requires specific IAM role configurations for each user, which could become complex to manage and might not provide the same level of straightforward access and editing capability for all users within the IAM group.

upvoted 1 times

tqiu654 1 year, 2 months ago

Selected Answer: A

Based on ChatGPT:A

upvoted 1 times

Jonalb 1 year, 3 months ago

Selected Answer: B

No AWS Lambda, você pode criar eventos de teste no console da AWS para invocar sua função e ver a resposta. Esses eventos de teste podem ser salvos e compartilhados com outros usuários IAM. Ao definir as configurações de compartilhamento de eventos para tornar os eventos de teste compartilháveis, você permite que todos os desenvolvedores do grupo de desenvolvedores IAM os usem e editem.

upvoted 1 times

DUBERS 1 year, 6 months ago

Would this not be C just because that's the only one that has the added security of the IAM roles?

upvoted 1 times

Cloud_Cloud 1 year, 9 months ago

Selected Answer: B

there is an option in lambda console to share the event with other users

upvoted 1 times

MrTee 1 year, 9 months ago

Selected Answer: A

I meant to select A

upvoted 3 times

MrTee 1 year, 9 months ago

Selected Answer: B

To create a set of test events that can be used by all developers in an IAM developer group and that are editable by any of the IAM users in the group, the company should create and store the test events in Amazon S3 as JSON objects and allow S3 bucket access to all IAM users (Option A). This will allow all developers in the IAM developer group to access and edit the test events as needed. The other options do not provide a way for multiple developers to access and edit the test events.

upvoted 1 times

Fyssy 1 year, 9 months ago

Selected Answer: C

Use roles. Not all IAM users

upvoted 1 times


A developer is configuring an application's deployment environment in AWS CodePipeline. The application code is stored in a GitHub repository. The developer wants to ensure that the repository package's unit tests run in the new deployment environment. The developer has already set the pipeline's source provider to GitHub and has specified the repository and branch to use in the deployment.

Which combination of steps should the developer take next to meet these requirements with the LEAST overhead? (Choose two.)

- A. Create an AWS CodeCommit project. Add the repository package's build and test commands to the project's buildspec.
- B. Create an AWS CodeBuild project. Add the repository package's build and test commands to the project's buildspec.
- C. Create an AWS CodeDeploy project. Add the repository package's build and test commands to the project's buildspec.
- D. Add an action to the source stage. Specify the newly created project as the action provider. Specify the build artifact as the action's input artifact.
- E. Add a new stage to the pipeline after the source stage. Add an action to the new stage. Specify the newly created project as the action provider. Specify the source artifact as the action's input artifact.

Correct Answer: BE

Community vote distribution



MrTee **Highly Voted** 1 year, 9 months ago

The correct answer is B and E

The buildspec file is a collection of build commands and related settings, in YAML format, that CodeBuild uses to run a build. By adding the build and test commands to the buildspec file, the developer can ensure that these commands are executed as part of the build process. Option E will ensure that the CodeBuild project is triggered as part of the pipeline and that the unit tests are run in the new deployment environment.

upvoted 21 times

imvb88 **Highly Voted** 1 year, 8 months ago

Selected Answer: BE

For those who just skim the question, keyword between D and E is "unit tests run in the new deployment environment.", which signifies a new stage should be created instead of just adding an action.

upvoted 14 times

sumanshu **Most Recent** 1 month, 1 week ago

Selected Answer: BE

AWS CodeBuild is specifically designed for building and testing code in CI/CD pipelines. The buildspec file is where the developer can specify commands to build and run unit tests for the application package.

A new stage after the source stage ensures that the unit tests run on the package retrieved from GitHub.

upvoted 2 times

tsangckl 7 months, 2 weeks ago

This appear at 17 Jun exam

upvoted 2 times

65703c1 8 months, 2 weeks ago

Selected Answer: BE

BE is the correct answer.

upvoted 1 times

SerialIDr 1 year ago

Selected Answer: BE

E. Add a new stage to the pipeline after the source stage: This is the correct step. The developer should add a new stage to the pipeline specifically for building and testing the code. Within this stage, an action should be added that specifies the AWS CodeBuild project (created in step B) as the action provider. The source artifact (code fetched from GitHub) should be specified as the action's input artifact.

So, the combination of steps that should be taken next to meet these requirements with the least overhead are:

B. Create an AWS CodeBuild project. Add the repository package's build and test commands to the project's buildspec.

E. Add a new stage to the pipeline after the source stage. Add an action to the new stage. Specify the newly created CodeBuild project as the action provider. Specify the source artifact as the action's input artifact.

upvoted 3 times

🗨️ 👤 **LR2023** 1 year, 1 month ago

Selected Answer: BD

Choosing D as that is the least overhead. There is already a stage and you need to add an action test

upvoted 1 times

🗨️ 👤 **LR2023** 1 year, 1 month ago

Sorry will go with BE after doing more research as unit tests cannot be run in source stage as an action

upvoted 2 times

🗨️ 👤 **marolisa** 1 year, 5 months ago

B e D.

https://docs.aws.amazon.com/pt_br/codebuild/latest/userguide/how-to-create-pipeline-add-test.html

upvoted 3 times

🗨️ 👤 **aaok** 1 year, 8 months ago

Selected Answer: BE

As MrTee says.

upvoted 3 times

An engineer created an A/B test of a new feature on an Amazon CloudWatch Evidently project. The engineer configured two variations of the feature (Variation A and Variation B) for the test. The engineer wants to work exclusively with Variation A. The engineer needs to make updates so that Variation A is the only variation that appears when the engineer hits the application's endpoint.

Which solution will meet this requirement?

- A. Add an override to the feature. Set the identifier of the override to the engineer's user ID. Set the variation to Variation A.
- B. Add an override to the feature. Set the identifier of the override to Variation A. Set the variation to 100%.
- C. Add an experiment to the project. Set the identifier of the experiment to Variation B. Set the variation to 0%.
- D. Add an experiment to the project. Set the identifier of the experiment to the AWS account's account ID. Set the variation to Variation A.

Correct Answer: A

Community vote distribution

A (100%)


 **Fyssy** Highly Voted 1 year, 9 months ago

Selected Answer: A

Overrides let you pre-define the variation for selected users. to always receive the editable variation.

<https://aws.amazon.com/blogs/aws/cloudwatch-evidently/>

upvoted 15 times

 **jipark** 1 year, 6 months ago

the key looks "override" and allow only "userID"

upvoted 1 times


 **Baba_Eni** Highly Voted 1 year, 7 months ago

Selected Answer: A

Check Bullet point 9 in the link below

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch-Evidently-newfeature.html>

upvoted 10 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: A

An override allows you to force specific variations of a feature for a subset of users based on identifiers like user ID. This is useful for testing or debugging without affecting other users or the experiment as a whole.

upvoted 2 times

 **nbxyzd** 3 months ago

Selected Answer: A

Quote: To specify that certain users always see a certain variation, choose Overrides, Add override. Then, specify a user by entering their user ID, account ID, or some other identifier in Identifier, and specify which variation they should see.

This can be useful for members of your own testing team or other internal users when you want to make sure they see a specific variation. The sessions of users who are assigned overrides do not contribute to launch or experiment metrics.

upvoted 1 times

 **michele740** 7 months, 2 weeks ago

Therefore, option A is the best choice to meet the requirement of ensuring the engineer exclusively sees Variation A.


upvoted 1 times

 **65703c1** 8 months, 2 weeks ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

 **badsati** 9 months, 3 weeks ago

Selected Answer: A

By adding an override to the feature and setting the identifier to the engineer's user ID, the engineer ensures that only their requests are directed to Variation A.

Setting the variation to Variation A explicitly assigns the desired variation to the engineer's requests, effectively ensuring they only experience Variation A.


Therefore, the correct solution is Option A.

upvoted 1 times

  **hsinchang** 1 year, 4 months ago

Set the variation to 0% or 100% makes no sense. Plus, the identifier should not be an account.

upvoted 2 times

  **ancomedian** 1 year, 6 months ago

Selected Answer: A

You have to give identifier

upvoted 1 times

A developer is working on an existing application that uses Amazon DynamoDB as its data store. The DynamoDB table has the following attributes: partNumber (partition key), vendor (sort key), description, productFamily, and productType. When the developer analyzes the usage patterns, the developer notices that there are application modules that frequently look for a list of products based on the productFamily and productType attributes.

The developer wants to make changes to the application to improve performance of the query operations.

Which solution will meet these requirements?

- A. Create a global secondary index (GSI) with productFamily as the partition key and productType as the sort key.
- B. Create a local secondary index (LSI) with productFamily as the partition key and productType as the sort key.
- C. Recreate the table. Add partNumber as the partition key and vendor as the sort key. During table creation, add a local secondary index (LSI) with productFamily as the partition key and productType as the sort key.
- D. Update the queries to use Scan operations with productFamily as the partition key and productType as the sort key.

Correct Answer: A

Community vote distribution

A (96%) 4%

 **Fyssy** Highly Voted 1 year, 3 months ago

Selected Answer: A

reate a Global Secondary Index (GSI): The developer should create a new GSI on the DynamoDB table with the productFamily attribute as the partition key and the productType attribute as the sort key. This will allow the application to perform fast queries on these attributes without scanning the entire table.


upvoted 10 times

 **Majong** Highly Voted 1 year, 2 months ago

Selected Answer: A

LSI can't be created on an already existing table and as Fyssy says. A - create new GSI will make the querying faster and you do not need to recreate the whole table.

upvoted 9 times


 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: A

A) Correct - GSI will give alternative Primary Key (HASH or HASH + Sort Key) from the base table.

B) Eliminated - An LSI shares the same partition key as the base table but allows a different sort key. LSIs must be defined at the time of table creation and cannot be added to an existing table.


upvoted 2 times

 **65703c1** 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

 **maurice2005** 4 months, 1 week ago

Selected Answer: C

Option C improve more performance. The question didn't care about effort

upvoted 1 times

 **SerialiDr** 6 months, 3 weeks ago

Selected Answer: A

This is a viable solution. A GSI allows you to query data using an alternate key, in this case, productFamily and productType. This would enable efficient queries based on these attributes, which is aligned with the observed usage patterns.

upvoted 2 times

  **winzzhhzzhh** 10 months, 4 weeks ago

Selected Answer: A

LSI: different sort key but the same partition key

GSI: different partition key and a different sort key

upvoted 5 times

A developer creates a VPC named VPC-A that has public and private subnets. The developer also creates an Amazon RDS database inside the private subnet of VPC-A. To perform some queries, the developer creates an AWS Lambda function in the default VPC. The Lambda function has code to access the RDS database. When the Lambda function runs, an error message indicates that the function cannot connect to the RDS database.

How can the developer solve this problem?

- A. Modify the RDS security group. Add a rule to allow traffic from all the ports from the VPC CIDR block.
- B. Redeploy the Lambda function in the same subnet as the RDS instance. Ensure that the RDS security group allows traffic from the Lambda function.
- C. Create a security group for the Lambda function. Add a new rule in the RDS security group to allow traffic from the new Lambda security group.
- D. Create an IAM role. Attach a policy that allows access to the RDS database. Attach the role to the Lambda function.

Correct Answer: B

Community vote distribution

B (73%)

C (27%)

 **MrTee** Highly Voted 1 year, 9 months ago

Selected Answer: B

To solve this problem, the developer should redeploy the Lambda function in the same subnet as the RDS instance and ensure that the RDS security group allows traffic from the Lambda function. This will allow the Lambda function to access the RDS database within the private subnet of VPC-A. The developer should also make sure that the Lambda function is configured with the appropriate network settings and permissions to access resources within the VPC.


upvoted 15 times

 **Fyssy** Highly Voted 1 year, 9 months ago

Selected Answer: B

Redeploy

upvoted 13 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B

Lambda function in the default VPC cannot communicate with the RDS in VPC-A, because they are in different VPCs with no connection (like a VPC peering or transit gateway).

upvoted 1 times

 **mallikarjun_angadi** 2 months, 2 weeks ago

Selected Answer: C

B is InCorrect because Deploying the Lambda function in the same private subnet as the RDS instance is not feasible, Lambda function needs to access other resources or services outside the VPC.

upvoted 2 times

 **mallikarjun_angadi** 2 months, 2 weeks ago

B is incorrect because deploying lambda in private subnet. Cannot access other resources

upvoted 1 times

 **raasankar** 4 months, 3 weeks ago

Selected Answer: B

C would need a vpc peering,

So B is the best option as we are redeploying to same subnet.

upvoted 1 times

 **65703c1** 8 months, 2 weeks ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗨️ **Vaibs099** 9 months, 1 week ago

they are in different VPC, hence C is not possible unless there is VPC Peering.

upvoted 1 times

🗨️ **badsati** 9 months, 3 weeks ago

Selected Answer: B

By deploying the Lambda function in the same subnet as the RDS instance (VPC-A), the Lambda function will have access to the resources within the same VPC, including the RDS database.

Additionally, the RDS security group should be configured to allow inbound traffic from the Lambda function's security group.

Therefore, the correct solution is Option B.

upvoted 2 times

🗨️ **DeaconStJohn** 10 months, 1 week ago

Selected Answer: B

Redeploy as no access has been set up between VPCs

upvoted 1 times

🗨️ **maurice2005** 10 months, 1 week ago

Selected Answer: B

A and B wont work since lambda is on default vpc which is not vpc-a

D won't work since since it's network access in the first place.

upvoted 1 times

🗨️ **SerialiDr** 11 months, 1 week ago

Selected Answer: B

Option B ("Redeploy the Lambda function in the same subnet as the RDS instance. Ensure that the RDS security group allows traffic from the Lambda function.") is the most accurate approach if the Lambda function and RDS are to communicate within the same VPC. It directly addresses the need for the Lambda function to access the VPC and the security group configuration.

upvoted 2 times

🗨️ **cauchy06** 1 year ago

Selected Answer: C

No need for redeploy. ChatGPT also says C.

upvoted 1 times

🗨️ **toan_nguyen** 11 months, 3 weeks ago

ChatGPT don't know anything. It's only read data

upvoted 5 times

🗨️ **SerialiDr** 1 year ago

Selected Answer: B

B. Redeploy the Lambda function in the same subnet as the RDS instance. Ensure that the RDS security group allows traffic from the Lambda function: This is a viable solution. Placing the Lambda function in the same VPC as the RDS instance (preferably in a private subnet for security reasons) and ensuring the security groups are correctly configured to allow traffic between the Lambda function and the RDS instance will enable connectivity.

C. Create a security group for the Lambda function. Add a new rule in the RDS security group to allow traffic from the new Lambda security group: This option would be correct if the Lambda function and the RDS instance were in the same VPC. However, since they are in different VPCs, simply adjusting security groups won't address the cross-VPC connectivity issue.

upvoted 6 times

🗨️ **nickolaj** 1 year ago

Selected Answer: B

Option C would be the correct choice, but it doesn't include the route configuration between subnets needed to access the RDS. I chose option B, but according to architectural best practices, it's not the ideal solution.

upvoted 2 times

🗨️ **a_win** 1 year, 1 month ago

Selected Answer: C

Seems more efficient solution.

upvoted 1 times

🗨️ 👤 **KarBiswa** 1 year, 1 month ago

Selected Answer: B

<https://docs.aws.amazon.com/vpc/latest/userguide/default-vpc.html>

The default VPC is the public subnet, this is the main trick

upvoted 1 times

🗨️ 👤 **BaYaga** 1 year, 1 month ago

Have you even read the documentation that you're providing. It says clearly:

You can use a default VPC as you would use any other VPC:

Add additional nondefault subnets.

Modify the main route table.

Add additional route tables.

Associate additional security groups.

Update the rules of the default security group.

Add AWS Site-to-Site VPN connections.

Add more IPv4 CIDR blocks.

Access VPCs in a remote Region by using a Direct Connect gateway. For information about Direct Connect gateway options, see Direct Connect gateways in the AWS Direct Connect User Guide.

You can use a default subnet as you would use any other subnet; add custom route tables and set network ACLs. You can also specify a specific default subnet when you launch an EC2 instance.

upvoted 1 times

A company runs an application on AWS. The company deployed the application on Amazon EC2 instances. The application stores data on Amazon Aurora.

The application recently logged multiple application-specific custom DECRYPT_ERROR errors to Amazon CloudWatch logs. The company did not detect the issue until the automated tests that run every 30 minutes failed. A developer must implement a solution that will monitor for the custom errors and alert a development team in real time when these errors occur in the production environment.

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

- A. Configure the application to create a custom metric and to push the metric to CloudWatch. Create an AWS CloudTrail alarm. Configure the CloudTrail alarm to use an Amazon Simple Notification Service (Amazon SNS) topic to send notifications.
- B. Create an AWS Lambda function to run every 5 minutes to scan the CloudWatch logs for the keyword DECRYPT_ERROR. Configure the Lambda function to use Amazon Simple Notification Service (Amazon SNS) to send a notification.
- C. Use Amazon CloudWatch Logs to create a metric filter that has a filter pattern for DECRYPT_ERROR. Create a CloudWatch alarm on this metric for a threshold ≥ 1 . Configure the alarm to send Amazon Simple Notification Service (Amazon SNS) notifications.
- D. Install the CloudWatch unified agent on the EC2 instance. Configure the application to generate a metric for the keyword DECRYPT_ERROR errors. Configure the agent to send Amazon Simple Notification Service (Amazon SNS) notifications.

Correct Answer: C

Community vote distribution

C (100%)

MrTee **Highly Voted** 1 year, 3 months ago

Selected Answer: C

To monitor for custom DECRYPT_ERROR errors and alert a development team in real time when these errors occur in the production environment with the least operational overhead, the developer should use Amazon CloudWatch Logs to create a metric filter that has a filter pattern for DECRYPT_ERROR. The developer should then create a CloudWatch alarm on this metric for a threshold ≥ 1 and configure the alarm to send Amazon Simple Notification Service (Amazon SNS) notifications (Option C). This solution will allow the developer to monitor for custom errors in real time and receive notifications when they occur with minimal operational overhead.

upvoted 9 times

sumanshu **Most Recent** 1 month, 1 week ago

Selected Answer: C

CloudWatch Logs Metric Filter allows you to directly track specific keywords (e.g., DECRYPT_ERROR) in logs and generate metrics without additional code or infrastructure. After creating a metric filter, you can configure a CloudWatch Alarm to monitor the metric and send an alert if the threshold is exceeded (e.g., ≥ 1 occurrence). Using Amazon SNS, the alarm can notify the development team in real time.

upvoted 1 times

sumanshu 1 month, 1 week ago

A) Eliminated - Custom metric creation by the application requires code changes in the application, which introduces development overhead.

B) Eliminated - Running a Lambda function periodically adds complexity, additional cost, and delays in detecting errors (up to 5 minutes).

D) Eliminated - installing and maintaining the CloudWatch agent introduces operational complexity

upvoted 1 times

65703c1 2 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

badsati 3 months, 3 weeks ago

Selected Answer: C

Options A, B, and D introduce additional complexity, operational overhead, and potential points of failure compared to Option C, which leverages native CloudWatch capabilities for log monitoring and alerting with minimal setup and maintenance overhead. Therefore, Option C is the most

suitable choice for meeting the requirements with the least operational overhead.

upvoted 1 times

🗨️ **SerialiDr** 6 months, 3 weeks ago

Selected Answer: C

This is a straightforward and effective solution. CloudWatch Logs allows you to create a metric filter for specific log patterns (such as DECRYPT_ERROR) and then create an alarm based on that metric. When the alarm is triggered, it can send a notification through Amazon SNS. This approach provides real-time monitoring with minimal operational overhead.

upvoted 1 times

🗨️ **hsinchang** 10 months, 3 weeks ago

Selected Answer: C

A and B are not real-time, and the CloudWatch unified agent in D is used to collect metrics and logs from EC2 instances and on-premises servers, not to send notifications.

So C.

upvoted 4 times

🗨️ **Fyssy** 1 year, 3 months ago

Selected Answer: C

CloudWatch Logs can use filter expressions. For example, find a specific IP inside of a log Or count occurrences of "ERROR" in your logs. Metric filters can be used to trigger CloudWatch alarms

upvoted 2 times

A developer created an AWS Lambda function that accesses resources in a VPC. The Lambda function polls an Amazon Simple Queue Service (Amazon SQS) queue for new messages through a VPC endpoint. Then the function calculates a rolling average of the numeric values that are contained in the messages. After initial tests of the Lambda function, the developer found that the value of the rolling average that the function returned was not accurate.


How can the developer ensure that the function calculates an accurate rolling average?

- A. Set the function's reserved concurrency to 1. Calculate the rolling average in the function. Store the calculated rolling average in Amazon ElastiCache.
- B. Modify the function to store the values in Amazon ElastiCache. When the function initializes, use the previous values from the cache to calculate the rolling average.
- C. Set the function's provisioned concurrency to 1. Calculate the rolling average in the function. Store the calculated rolling average in Amazon ElastiCache.
- D. Modify the function to store the values in the function's layers. When the function initializes, use the previously stored values to calculate the rolling average.

Correct Answer: A


Community vote distribution

A (57%) B (42%)

 **eboehm** Highly Voted 1 year, 7 months ago

Selected Answer: A


You need to set the reserved concurrency to 1 otherwise multiple functions could run at the same time causing the math to be off. Also there was a similar question in another practice exam set that stated the same thing
upvoted 25 times

 **jipark** 1 year, 6 months ago
reserve concurrency 1 means poll in order.
this looks answer.
upvoted 2 times

 **MrTee** Highly Voted 1 year, 9 months ago

Selected Answer: B

By using ElastiCache, the Lambda function can store the values of the previous messages it has received, which can be used to calculate an accurate rolling average.
upvoted 18 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: A

B) Eliminated - If multiple instances of the Lambda function run concurrently, they could retrieve or modify the same data in ElastiCache, leading to inaccuracies in the rolling average.

C) Eliminated - Provisioned concurrency does not prevent concurrency issues


D) Eliminated - Lambda layers are used to share common code or dependencies across Lambda functions but are not designed for storing or persisting data.
upvoted 1 times

 **ShakthiGCP** 1 month, 1 week ago

Selected Answer: B

Reserved concurrency: Setting reserved concurrency to 1 ensures only one instance of the function executes at a time. While this might prevent race conditions, it doesn't address the core issue of calculating the rolling average across multiple Lambda invocations.

u
upvoted 1 times

 **trieudo** 1 month, 2 weeks ago

Selected Answer: A

==> Discard B: There are not limit concurrency in elastic cache, multiple lambda can access make race condition ==> wrong output
upvoted 1 times

🗨️ **AmitRanchi** 3 months, 3 weeks ago

Selected Answer: A

Reserved Concurrency of 1 ensures that only one instance of the Lambda function runs at a time, which is crucial for calculating an accurate rolling average. If multiple instances of the Lambda function run simultaneously (without concurrency control), each instance might have incomplete or inconsistent data, leading to incorrect results for the rolling average.

Storing the rolling average in Amazon ElastiCache provides a persistent, in-memory storage that the function can access quickly during each invocation. ElastiCache can store the current state of the rolling average and any relevant historical data, ensuring that each invocation of the Lambda function has access to the most recent values for calculation

upvoted 1 times

🗨️ **Saurabh04** 5 months, 4 weeks ago

Correct answer should be D. Storing values in the function's layers and using them during initialization is straightforward. This Approach avoids external services like Elasticache.

upvoted 1 times

🗨️ **rue_** 3 months ago

layers are not meant to store dynamic data. it's for sharing libraries and static code across lambda functions

upvoted 1 times

🗨️ **lozou** 6 months, 3 weeks ago

Selected Answer: B

elasticache for keeping the values from previous invocation

upvoted 1 times

🗨️ **guidosolano** 7 months ago

Selected Answer: B

A y C limitan mucho la capacidad de Lamda. Voy con B

upvoted 1 times

🗨️ **65703c1** 8 months, 2 weeks ago

Selected Answer: A

A is the correct answer.

upvoted 2 times

🗨️ **SathyaJS** 10 months, 1 week ago

Selected Answer: B

As per Gemini below is why A is incorrect

A. Reserved concurrency: Setting reserved concurrency to 1 ensures only one instance of the function executes at a time. While this might prevent race conditions, it doesn't address the core issue of calculating the rolling average across multiple Lambda invocations.

upvoted 1 times

🗨️ **yingying920928** 10 months, 2 weeks ago

Selected Answer: B

While limiting concurrency (A or C) can help manage the function's execution rate and scale and prevent throttling, it is not directly related to ensuring the accuracy of calculating rolling average. Instead, focusing on proper state management and data consistency mechanisms (using ElastiCache) is key to achieving accurate results in this scenario.

upvoted 2 times

🗨️ **maurice2005** 10 months, 1 week ago

that's the correct way in real world but nothing about state management or data consistency is mentioned in B. But A has it although it's not the good real world to do so.

upvoted 1 times

🗨️ **SerialiDr** 11 months, 1 week ago

Selected Answer: A

Option A ("Set the function's reserved concurrency to 1. Calculate the rolling average in the function. Store the calculated rolling average in Amazon ElastiCache.") is the most suitable solution. It ensures that only one instance of the Lambda function processes messages at any given time, maintaining the sequence of message processing which is crucial for an accurate rolling average calculation. Additionally, using Amazon

ElastiCache to store and retrieve the rolling average across invocations addresses the statelessness of AWS Lambda, enabling stateful processing.

upvoted 4 times

🗨️ **d323bvmiqj** 11 months, 2 weeks ago

Selected Answer: A

What if one of the instances freezes and holds one of the values for some time, not updating cache, while the others continue calculating the avg giving wrong output ?

upvoted 3 times

🗨️ **SerialiDr** 1 year ago

Selected Answer: B

By storing individual message values in ElastiCache (a fast, in-memory data store), the Lambda function can retrieve these values upon initialization to calculate an accurate rolling average. This approach effectively maintains state across Lambda invocations.

upvoted 2 times

🗨️ **Chimzi** 1 year ago

Selected Answer: B

Using ElastiCache allows you to maintain a shared state across all instances of your Lambda function

upvoted 1 times

🗨️ **ShinobiGrappler** 1 year, 1 month ago

Selected Answer: A

This approach controls concurrency by ensuring only one instance runs at a time. Provisioned concurrency also has the added benefit of reducing cold start latency. Storing the rolling average in ElastiCache is a good practice for maintaining state. However, like option A, it may limit the function's throughput.

upvoted 4 times

A developer is writing unit tests for a new application that will be deployed on AWS. The developer wants to validate all pull requests with unit tests and merge the code with the main branch only when all tests pass.

The developer stores the code in AWS CodeCommit and sets up AWS CodeBuild to run the unit tests. The developer creates an AWS Lambda function to start the CodeBuild task. The developer needs to identify the CodeCommit events in an Amazon EventBridge event that can invoke the Lambda function when a pull request is created or updated.

Which CodeCommit event will meet these requirements?

- A.

```
{
  "source": ["aws.codecommit"],
  "detail": {
    "event": ["pullRequestMergeStatusUpdated"],
  }
}
```
- B.

```
{
  "source": ["aws.codecommit"],
  "detail": {
    "event": ["pullRequestApprovalRuleCreated"]
  }
}
```
- C.

```
{
  "source": ["aws.codecommit"],
  "detail": {
    "event": ["pullRequestSourceBranchUpdated", "pullRequestCreated"]
  }
}
```
- D.

```
{
  "source": ["aws.codecommit"],
  "detail": {
    "event": ["pullRequestUpdated", "pullRequestSourceBranchCreated"]
  }
}
```

Correct Answer: C

Community vote distribution

C (86%)


14%

 **csG13** Highly Voted 1 year, 1 month ago

Selected Answer: C

It's definitely C. Events in answer D are not real. A & B are clearly wrong since two events are required.


upvoted 8 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: C

Using both events together ensures that the Lambda function is invoked for both new pull requests and updates to existing pull requests.

upvoted 1 times

 **65703c1** 2 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

 **Vaibs099** 3 months, 1 week ago

<https://aws.amazon.com/blogs/devops/automated-code-review-on-pull-requests-using-aws-codecommit-and-aws-codebuild/>

upvoted 2 times

 **Dushank** 10 months, 4 weeks ago

Answer is C. There's no event call pullRequestUpdated

upvoted 4 times

 **Majong** 1 year, 2 months ago

Selected Answer: C

Two events is needed so A and B is no.

The events mentioned in D does not exist as Zodraz says (just look in the link)

upvoted 4 times

🗨️ 👤 **Prem28** 1 year, 2 months ago

Selected Answer: C

its c ,Event mentioned in D not Exist

upvoted 3 times

🗨️ 👤 **zodraz** 1 year, 2 months ago

Selected Answer: C

It' s C. Any of the events mentioned on D exist. <https://docs.aws.amazon.com/codecommit/latest/userguide/monitoring-events.html#pullRequestSourceBranchUpdated>

upvoted 3 times

🗨️ 👤 **zodraz** 1 year, 2 months ago

It' s C. Any of the events mentioned on D exist. <https://docs.aws.amazon.com/codecommit/latest/userguide/monitoring-events.html#pullRequestSourceBranchUpdated>

upvoted 2 times

🗨️ 👤 **Fyssy** 1 year, 3 months ago

Selected Answer: D

```
"detail": {
```

```
"event": ["pullRequestCreated", "pullRequestSourceBranchUpdated"]
```

upvoted 3 times

🗨️ 👤 **Arktos** 2 months, 1 week ago

you create a pull request and update a branch so the answer is C

upvoted 1 times

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

Correct Answer: A

Community vote distribution

A (100%)

🗨️ 👤 **Naj_64** Highly Voted 👍 1 year ago

Selected Answer: A

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instancedata-data-retrieval.html>

upvoted 6 times

🗨️ 👤 **SerialiDr** Highly Voted 👍 6 months, 3 weeks ago

Selected Answer: A

This is the correct approach. The instance metadata includes details such as the instance's public IPv4 address. The application can make a request to this URL, specifically to <http://169.254.169.254/latest/meta-data/public-ipv4>, to retrieve the public IPv4 address of the instance.

upvoted 6 times

🗨️ 👤 **sumanshu** Most Recent 🕒 1 month, 1 week ago

Selected Answer: A

The EC2 instance metadata service provides details about the instance, such as its public and private IP addresses. To find the public IPv4 address, you can query the metadata

upvoted 1 times

🗨️ 👤 **65703c1** 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 2 times

🗨️ 👤 **zodraz** 1 year, 2 months ago

Selected Answer: A

You can retrieve ip through <http://169.254.169.254/latest/meta-data/local-ipv4> or <http://169.254.169.254/latest/meta-data/public-ipv4>

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instancedata-data-retrieval.html>

upvoted 2 times

🗨️ 👤 **zodraz** 1 year, 2 months ago

Selected Answer: A

It' s C. Any of the events mentioned on D exist. <https://docs.aws.amazon.com/codecommit/latest/userguide/monitoring-events.html#pullRequestSourceBranchUpdated>

upvoted 2 times

🗨️ 👤 **zodraz** 1 year, 2 months ago

Please remove this comment @admin

upvoted 4 times

An application under development is required to store hundreds of video files. The data must be encrypted within the application prior to storage, with a unique key for each video file.

How should the developer code the application?

- A. Use the KMS Encrypt API to encrypt the data. Store the encrypted data key and data.
- B. Use a cryptography library to generate an encryption key for the application. Use the encryption key to encrypt the data. Store the encrypted data.
- C. Use the KMS GenerateDataKey API to get a data key. Encrypt the data with the data key. Store the encrypted data key and data.
- D. Upload the data to an S3 bucket using server side-encryption with an AWS KMS key.

Correct Answer: C

Community vote distribution

C (100%)

MrTee **Highly Voted** 1 year, 3 months ago

Selected Answer: C

option C: use the KMS GenerateDataKey API to get a data key. Encrypt the data with the data key. Store the encrypted data key and data.
upvoted 12 times

sumanshu **Most Recent** 1 month, 1 week ago

Selected Answer: C

A) Eliminated - The Encrypt API in AWS KMS is designed for small amounts of data (4 KB)

For large files, it is more efficient and scalable to use data keys (using the GenerateDataKey API)
upvoted 2 times

sumanshu 1 month, 1 week ago

D) Eliminated - this approach does not meet the requirement to encrypt data within the application prior to storage,
upvoted 2 times

65703c1 2 months, 1 week ago

Selected Answer: C

C is the correct answer.
upvoted 1 times

SerialiDr 6 months, 3 weeks ago

Selected Answer: C

This is the most suitable option. AWS KMS's GenerateDataKey API provides a unique data key for each invocation, which can be used to encrypt each video file. The data key itself is also returned in an encrypted form, which can be safely stored alongside the encrypted data. This approach satisfies the requirement of unique encryption for each file and securely manages the encryption keys.
upvoted 3 times

Tinez 9 months, 1 week ago

Option C seems correct
upvoted 1 times

hsinchang 10 months, 3 weeks ago

Selected Answer: C

A and B cannot meet the requirement of having a unique key for each file, and D cannot meet the requirement of encrypting the data within the application.
C meets all requirements.
upvoted 2 times

A company is planning to deploy an application on AWS behind an Elastic Load Balancer. The application uses an HTTP/HTTPS listener and must access the client IP addresses.

Which load-balancing solution meets these requirements?

- A. Use an Application Load Balancer and the X-Forwarded-For headers.
- B. Use a Network Load Balancer (NLB). Enable proxy protocol support on the NLB and the target application.
- C. Use an Application Load Balancer. Register the targets by the instance ID.
- D. Use a Network Load Balancer and the X-Forwarded-For headers.

Correct Answer: A

Community vote distribution

A (100%)

MrTee **Highly Voted** 1 year, 3 months ago

Selected Answer: A

Use an Application Load Balancer (ALB) and the X-Forwarded-For headers. When an ALB is used, the X-Forwarded-For header can be used to pass the client IP address to the backend servers.

upvoted 10 times

sumanshu **Most Recent** 1 month, 1 week ago

Selected Answer: A

Application Load Balancer (ALB) is a Layer 7 (HTTP/HTTPS) load balancer that supports handling HTTP headers, including the X-Forwarded-For header.

B/D - Eliminated - Network Load Balancer (NLB) operates at Layer 4 (TCP) and does not inherently handle HTTP headers

C) Eliminated - it does not address the requirement of accessing client IP addresses.

upvoted 2 times

65703c1 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

topicsquestions 3 months, 2 weeks ago

Selected Answer: A

Option a

upvoted 1 times

SerialiDr 6 months, 3 weeks ago

Selected Answer: A

An Application Load Balancer (ALB) operates at the application layer (Layer 7) of the OSI model and supports HTTP/HTTPS traffic. It adds the X-Forwarded-For header to the request as it forwards it to the target, which contains the original client's IP address. This allows the application behind the ALB to access the client IP addresses.

upvoted 4 times

HuiHsin 1 year, 1 month ago

is A

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/x-forwarded-headers.html>

<https://aws.amazon.com/elasticloadbalancing/features/?nc=sn&loc=2>

upvoted 4 times

A developer wants to debug an application by searching and filtering log data. The application logs are stored in Amazon CloudWatch Logs. The developer creates a new metric filter to count exceptions in the application logs. However, no results are returned from the logs.

What is the reason that no filtered results are being returned?

- A. A setup of the Amazon CloudWatch interface VPC endpoint is required for filtering the CloudWatch Logs in the VPC.
- B. CloudWatch Logs only publishes metric data for events that happen after the filter is created.
- C. The log group for CloudWatch Logs should be first streamed to Amazon OpenSearch Service before metric filtering returns the results.
- D. Metric data points for logs groups can be filtered only after they are exported to an Amazon S3 bucket.

Correct Answer: B

Community vote distribution

B (100%)

🗉 **zodraz** Highly Voted 1 year, 2 months ago

Selected Answer: B

Filters do not retroactively filter data. Filters only publish the metric data points for events that happen after the filter was created.

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/MonitoringLogData.html>

upvoted 14 times

🗉 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: B

Metric filters in Amazon CloudWatch Logs only work on new logs generated after the filter is created. If the filter was created, but no logs were generated afterward (or if you're trying to filter older log entries), no metric data points would be produced.

upvoted 1 times

🗉 **65703c1** 2 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗉 **SerialiDr** 6 months, 3 weeks ago

Selected Answer: B

CloudWatch Logs metric filters apply to new log events only after the filter is created. They do not retroactively analyze or filter log events that were ingested before the creation of the metric filter. Therefore, if the log events in question were ingested before the metric filter was created, they would not trigger the filter or generate metric data.

upvoted 2 times

🗉 **Dushank** 10 months, 3 weeks ago

Selected Answer: B

Metric filters in Amazon CloudWatch Logs are applied only to new log events. If you create a metric filter and are looking to count exceptions, the filter will only apply to log events generated after the metric filter was created. Existing logs are not scanned.

upvoted 4 times

A company is planning to use AWS CodeDeploy to deploy an application to Amazon Elastic Container Service (Amazon ECS). During the deployment of a new version of the application, the company initially must expose only 10% of live traffic to the new version of the deployed application. Then, after 15 minutes elapse, the company must route all the remaining live traffic to the new version of the deployed application.

Which CodeDeploy predefined configuration will meet these requirements?

- A. CodeDeployDefault.ECSCanary10Percent15Minutes
- B. CodeDeployDefault.LambdaCanary10Percent5Minutes
- C. CodeDeployDefault.LambdaCanary10Percent15Minutes
- D. CodeDeployDefault.ECSLinear10PercentEvery1Minutes

Correct Answer: A

Community vote distribution

A (100%)

 **zodraz** Highly Voted 1 year, 2 months ago

Selected Answer: A

<https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-configurations.html>

upvoted 9 times

 **sumanshu** Most Recent 1 month, 1 week ago

Selected Answer: A

B/C - Eliminated - This configuration is for Lambda deployments, not ECS deployments

D) Eliminated - This configuration specifies a linear deployment, meaning that the traffic will be gradually shifted to the new version by 10% every 1 minute


upvoted 1 times

 **65703c1** 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

 **SerialiDr** 6 months, 3 weeks ago

Selected Answer: A

This configuration aligns with the company's requirement. It specifies a "canary" deployment where initially only 10% of live traffic is exposed to the new version of the application. After a period of 15 minutes, the remaining 90% of the traffic is shifted to the new version. This approach allows for monitoring the new version with a small portion of traffic before fully deploying it.

upvoted 4 times

 **Dushank** 10 months, 3 weeks ago

Selected Answer: A

This predefined deployment configuration for AWS CodeDeploy with Amazon ECS will initially shift 10% of the traffic to the new version and wait for 15 minutes before shifting the remaining 90% of the traffic to the new version.

upvoted 4 times

A company hosts a batch processing application on AWS Elastic Beanstalk with instances that run the most recent version of Amazon Linux. The application sorts and processes large datasets.

In recent weeks, the application's performance has decreased significantly during a peak period for traffic. A developer suspects that the application issues are related to the memory usage. The developer checks the Elastic Beanstalk console and notices that memory usage is not being tracked.

How should the developer gather more information about the application performance issues?

- A. Configure the Amazon CloudWatch agent to push logs to Amazon CloudWatch Logs by using port 443.
- B. Configure the Elastic Beanstalk `.ebextensions` directory to track the memory usage of the instances.
- C. Configure the Amazon CloudWatch agent to track the memory usage of the instances.
- D. Configure an Amazon CloudWatch dashboard to track the memory usage of the instances.

Correct Answer: C

Community vote distribution

C (60%)

B (40%)

 **MrTee** Highly Voted 1 year, 9 months ago

Selected Answer: C

Configure the Amazon CloudWatch agent to track the memory usage of the instances.
upvoted 16 times

 **xdkonorek2** 1 year, 1 month ago

Using elastic beanstalk `.ebextensions` dir
upvoted 1 times


 **eboehm** Highly Voted 1 year, 7 months ago

Selected Answer: B

for elastic beanstalk you make this configuration in the `.ebextensions` folder
<https://repost.aws/knowledge-center/elastic-beanstalk-memory-metrics-windows>
upvoted 15 times

 **DumPisach** 1 year, 7 months ago

But the question says Linux
upvoted 3 times

 **Naj_64** 1 year, 6 months ago

Applies to Linux as well:
<https://medium.com/tomincode/cloudwatch-memory-monitoring-for-elastic-beanstalk-1caa98d57d5c>
upvoted 1 times

 **maurice2005** 10 months, 2 weeks ago

`ebextensions` is deprecated for linux newest version which is mentioned in the question
upvoted 4 times

 **sumanshu** Most Recent 1 month ago

Selected Answer: C

A) Eliminated - Logs provide insights into application behavior but won't help in monitoring resource usage such as memory.

B) Eliminated - The `.ebextensions` directory is used for customizing Elastic Beanstalk environments, such as installing software or modifying instance configurations.

C) Correct - The Amazon CloudWatch agent can collect system-level metrics like memory usage, which is not tracked by default in CloudWatch for EC2 instances

D) Eliminated - A CloudWatch dashboard is used to visualize metrics that are already being tracked.

upvoted 2 times

🗨️ **ShakthiGCP** 1 month, 1 week ago

Selected Answer: C

.ebextensions is for customizing your Elastic Beanstalk environment. While you can use it to install monitoring tools, directly configuring it to track memory isn't the standard or most efficient approach. So Answer is 'C'

upvoted 1 times

🗨️ **Saudis** 4 months, 3 weeks ago

Selected Answer: C

CloudWatch agent track the memory and any types of system component (hardware)

upvoted 1 times

🗨️ **65703c1** 8 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗨️ **SathyaJS** 10 months, 1 week ago

Answer is : C

from Gemini here is why B is incorrect

B. .ebextensions directory: While .ebextensions can be used for configuration within Elastic Beanstalk, it's not designed to track memory usage directly.

upvoted 2 times

🗨️ **SerialiDr** 11 months, 1 week ago

Selected Answer: C

This option allows the developer to gather detailed performance metrics, including memory usage, from the EC2 instances. By configuring the CloudWatch agent, the developer can monitor the memory usage in real-time and analyze historical data to identify trends or patterns that may be affecting the application's performance. This approach provides actionable insights with minimal overhead and without the need for custom logging or external tools.

upvoted 4 times

🗨️ **prathameshpathak** 1 year ago

Selected Answer: C

.....

upvoted 1 times

🗨️ **SerialiDr** 1 year ago

Selected Answer: C

This is the most direct and appropriate solution. By installing and configuring the Amazon CloudWatch agent on the Elastic Beanstalk instances, the developer can collect detailed system-level metrics, such as memory usage, and send them to CloudWatch for monitoring and analysis.

upvoted 1 times

🗨️ **Chimzi** 1 year ago

Selected Answer: C

The .ebextensions directory is used for customizing the environment (installing packages, running scripts...) it can't track memory usage alone.

upvoted 2 times

🗨️ **Chimzi** 1 year ago

Selected Answer: C

No Discussion

upvoted 1 times

🗨️ **JohnPI** 1 year ago

Selected Answer: C

We configure the agent not the directory itself.

upvoted 1 times

🗨️ **a_win** 1 year, 1 month ago

Selected Answer: C

The Amazon CloudWatch agent can be configured to collect various metrics, including memory usage, from the instances. By setting up the CloudWatch agent to monitor memory metrics, the developer can get insights into the memory usage patterns during peak traffic periods. This

data can help diagnose if memory constraints are causing the performance degradation.

upvoted 1 times

🗨️ **[Removed]** 1 year, 1 month ago

Selected Answer: B

it should be B:

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/customize-containers-cw.html#customize-containers-cw-update-roles>

upvoted 3 times

🗨️ **KarBiswa** 1 year, 1 month ago

Selected Answer: C

Going with C after going through this link:

<https://repost.aws/knowledge-center/elastic-beanstalk-memory-cpu-issues>

upvoted 1 times

🗨️ **TallManDan** 1 year, 2 months ago

It requires both B and C. I'm guessing the question is supposed to say "Select Two".

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/customize-containers-cw.html>

upvoted 1 times

A developer is building a highly secure healthcare application using serverless components. This application requires writing temporary data to /tmp storage on an AWS Lambda function.

How should the developer encrypt this data?

- A. Enable Amazon EBS volume encryption with an AWS KMS key in the Lambda function configuration so that all storage attached to the Lambda function is encrypted.
- B. Set up the Lambda function with a role and key policy to access an AWS KMS key. Use the key to generate a data key used to encrypt all data prior to writing to /tmp storage.
- C. Use OpenSSL to generate a symmetric encryption key on Lambda startup. Use this key to encrypt the data prior to writing to /tmp.
- D. Use an on-premises hardware security module (HSM) to generate keys, where the Lambda function requests a data key from the HSM and uses that to encrypt data on all requests to the function.

Correct Answer: B

Community vote distribution

B (100%)


 **sumanshu** 1 month ago

Selected Answer: B

A) Eliminated - AWS Lambda does not use Amazon EBS volumes for its temporary /tmp storage.

C) Eliminated - OpenSSL could theoretically generate a symmetric key, this approach is less secure and more error-prone compared to using AWS KMS


D) Eliminated - Using an on-premises HSM adds unnecessary complexity and latency to a serverless application
upvoted 1 times

 **albert_kuo** 4 months, 1 week ago

Selected Answer: B

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms:Encrypt",
        "kms:Decrypt",
        "kms:GenerateDataKey",
        "kms:GenerateDataKeyWithoutPlaintext",
        "kms:ReEncrypt*"
      ],
      "Resource": "arn:aws:kms:region:account-id:key/key-id"
    }
  ]
}
```

upvoted 2 times

 **65703c1** 8 months, 1 week ago

Selected Answer: B

B is the correct answer.
upvoted 1 times

 **SerialiDr** 1 year ago

Selected Answer: B

AWS Key Management Service (KMS) provides secure management of encryption keys. The Lambda function can use a KMS key to generate data keys for encrypting and decrypting data. The Lambda function would require appropriate permissions to access the KMS key. This approach provides a high level of security, which is essential for a healthcare application.

upvoted 4 times

🗨️ **Milan61** 1 year, 3 months ago

B is the solution

upvoted 1 times

🗨️ **Yuxing_Li** 1 year, 5 months ago

Selected Answer: B

Go with B

upvoted 2 times

🗨️ **abdelbz16** 1 year, 9 months ago

Selected Answer: B

B is the best solution

upvoted 4 times

🗨️ **MrTee** 1 year, 9 months ago

Selected Answer: B

is the best solution for encrypting temporary data written to /tmp storage on an AWS Lambda function

upvoted 4 times


A developer has created an AWS Lambda function to provide notification through Amazon Simple Notification Service (Amazon SNS) whenever a file is uploaded to Amazon S3 that is larger than 50 MB. The developer has deployed and tested the Lambda function by using the CLI. However, when the event notification is added to the S3 bucket and a 3,000 MB file is uploaded, the Lambda function does not launch.

Which of the following is a possible reason for the Lambda function's inability to launch?

- A. The S3 event notification does not activate for files that are larger than 1,000 MB.
- B. The resource-based policy for the Lambda function does not have the required permissions to be invoked by Amazon S3.
- C. Lambda functions cannot be invoked directly from an S3 event.
- D. The S3 bucket needs to be made public.

Correct Answer: B

Community vote distribution



Jamshif01 Highly Voted 1 year, 2 months ago

Selected Answer: B

B - is right answer

A is incorrect because the size of the file should not affect whether the event notification is triggered.

C is incorrect because Lambda functions can indeed be invoked directly from an S3 event.

D is incorrect because the S3 bucket does not need to be made public for the Lambda function to be invoked.

(c)chatgpt

upvoted 12 times

sumanshu Most Recent 1 month ago

Selected Answer: B

A) Eliminated - S3 event notifications do not have a size limitation for triggering Lambda functions.

C) Eliminated - AWS Lambda functions can be triggered directly by S3 event notifications.

D) Eliminated - Making the S3 bucket public is not required for it to trigger an event notification

upvoted 2 times

65703c1 2 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

ec8or 4 months, 1 week ago

Answer is A. The notification work normally. They should work for files up to 5TB but not once the dev uploads a file of 3000MB. Seems the dev did not set up the notification correctly for files of this size.

upvoted 1 times

Melisa202401 5 months, 2 weeks ago

Why answer B while dev deployed and tested via CLI is ok, but the reason would be lack of resource policy?

upvoted 1 times

TheFivePips 4 months, 4 weeks ago

The Lambda function, which reacts to an event and sends a message, can be invoked manually using the AWS CLI or other methods, even if there are issues with the S3 event triggering mechanism. This is because manual invocation typically bypasses the event source (S3 in this case) and directly triggers the Lambda function.

While Lambda functions can be invoked manually, they also need the correct permissions to be triggered by specific event sources like S3. If the resource-based policy for the Lambda function does not have the necessary permissions, it may not be invoked by S3 events.

upvoted 1 times

🗨️ 👤 **Prem28** 1 year, 2 months ago

B

A. The S3 event notification does not activate for files that are larger than 1,000 MB. This is not the case. S3 event notifications can activate for files that are larger than 1,000 MB.

C. Lambda functions cannot be invoked directly from an S3 event. This is also not the case. Lambda functions can be invoked directly from an S3 event.

D. The S3 bucket needs to be made public. This is not necessary. The S3 bucket does not need to be made public in order for the Lambda function to be invoked.

upvoted 2 times

🗨️ 👤 **chumji** 1 year, 2 months ago

Selected Answer: B

answer is B

upvoted 2 times

🗨️ 👤 **junrun3** 1 year, 2 months ago

Selected Answer: A

Answer A

upvoted 1 times

🗨️ 👤 **junrun3** 1 year, 2 months ago

not A, answer is B

upvoted 4 times

A developer is creating a Ruby application and needs to automate the deployment, scaling, and management of an environment without requiring knowledge of the underlying infrastructure.

Which service would best accomplish this task?

- A. AWS CodeDeploy
- B. AWS CloudFormation
- C. AWS OpsWorks
- D. AWS Elastic Beanstalk

Correct Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **Prem28** Highly Voted 👍 1 year, 8 months ago

answer- d

AWS CodeDeploy can automate the deployment of code to any instance, including Amazon EC2 instances and on-premises servers. However, it does not provide the same level of automation as Elastic Beanstalk, and it requires more manual intervention from developers.

AWS CloudFormation can help you model and set up your AWS resources. However, it does not provide any automation for deploying or managing applications.

AWS OpsWorks is a configuration management service that provides managed instances of Chef and Puppet. However, it is not as easy to use as Elastic Beanstalk, and it does not provide the same level of automation for deploying or managing applications.

upvoted 14 times

🗳️ 👤 **zodraz** Highly Voted 👍 1 year, 8 months ago

Selected Answer: D

<https://www.examttopics.com/discussions/amazon/view/88659-exam-aws-certified-developer-associate-topic-1-question-197/>

upvoted 6 times

🗳️ 👤 **sumanshu** Most Recent 🕒 1 month ago

Selected Answer: D

A) Eliminated - It does not handle the environment's scaling or management

B) Eliminated - AWS CloudFormation is used to define and provision infrastructure as code

C) Eliminated -AWS OpsWorks is a configuration management service that supports Chef and Puppet to automate deployments and configurations

D) Correct

upvoted 1 times

🗳️ 👤 **nbxyzd** 3 months ago

Selected Answer: D

without requiring knowledge of the underlying infrastructure ---> AWS Elastic Beanstalk

upvoted 1 times

🗳️ 👤 **65703c1** 8 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

🗳️ 👤 **Dushank** 1 year, 4 months ago

Selected Answer: D

AWS Elastic Beanstalk is designed for developers like the one in your scenario who want to deploy and manage applications without worrying about the underlying infrastructure. It automates the deployment process and automatically handles capacity provisioning, load balancing, auto-scaling, and application health monitoring. You can use it with various platforms including Ruby.

upvoted 3 times

A company has a web application that is deployed on AWS. The application uses an Amazon API Gateway API and an AWS Lambda function as its backend.

The application recently demonstrated unexpected behavior. A developer examines the Lambda function code, finds an error, and modifies the code to resolve the problem. Before deploying the change to production, the developer needs to run tests to validate that the application operates properly.


The application has only a production environment available. The developer must create a new development environment to test the code changes. The developer must also prevent other developers from overwriting these changes during the test cycle.

Which combination of steps will meet these requirements with the LEAST development effort? (Choose two.)

- A. Create a new resource in the current stage. Create a new method with Lambda proxy integration. Select the Lambda function. Add the hotfix alias. Redeploy the current stage. Test the backend.
- B. Update the Lambda function in the API Gateway API integration request to use the hotfix alias. Deploy the API Gateway API to a new stage named hotfix. Test the backend.
- C. Modify the Lambda function by fixing the code. Test the Lambda function. Create the alias hotfix. Point the alias to the \$LATEST version.
- D. Modify the Lambda function by fixing the code. Test the Lambda function. When the Lambda function is working as expected, publish the Lambda function as a new version. Create the alias hotfix. Point the alias to the new version.
- E. Create a new API Gateway API for the development environment. Add a resource and method with Lambda integration. Choose the Lambda function and the hotfix alias. Deploy to a new stage. Test the backend.

Correct Answer: BD

Community vote distribution



🗨️ 👤 **sumanshu** 1 month ago

Selected Answer: BD

A) Eliminated - If we create a new resource (e.g., an API path like /new-feature) The current stage (prod) is being modified, which is risky because the production environment is shared.

C) Eliminated - \$LATEST is not stable. It is always the most recent version of the code, and other developers can overwrite it with new changes during the testing phase.

E) Eliminated - Creating a completely new API Gateway API requires significant additional effort, upvoted 1 times

🗨️ 👤 **65703c1** 2 months, 1 week ago

Selected Answer: BD

BD is the correct answer.

upvoted 1 times

🗨️ 👤 **KillThemWithKindness** 5 months, 1 week ago

Selected Answer: BD

Not C, you can't use an unqualified ARN (\$LATEST) to create an alias.

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-versions.html>

E

After the initial deployment, you can add more stages and associate them with existing deployments. You can use the API Gateway console to create a new stage, or you can choose an existing stage while deploying an API. You can add a new stage to an API deployment before redeploying the API.

<https://docs.aws.amazon.com/apigateway/latest/developerguide/stages.html>

upvoted 2 times

🗨️ **CrescentShared** 5 months, 3 weeks ago

Selected Answer: BE

I cannot find another choice that meets this requirement.

"The developer must create a new development environment to test the code changes."

upvoted 1 times

🗨️ **SerialiDr** 6 months, 3 weeks ago

Selected Answer: BD

The order is D and then B

upvoted 2 times

🗨️ **Ponyi** 8 months, 4 weeks ago

Selected Answer: BD

Why D over C?

Versions are immutable. \$Latest is mutable, which means anyone access to Lambda can edit and deploy a new code. The question simply doesn't want that.

Why B over E?

You don't need to create a whole new API to test some new feature. You can simply achieve this by deploying it to a different stage. Afterwards, you can redirect the users to a new stage or do A/B testing.

upvoted 3 times

🗨️ **r3mo** 1 year ago

C - D.

C vs B : option C is preferred over option B because it provides a more isolated and controlled environment for testing the hotfix without directly affecting the production environment. It gives you the flexibility to iterate on the hotfix if needed and promotes a safer development and testing process.

D vs E : Option E is preferred over option D because it provides a more isolated and controlled environment for testing the hotfix. It avoids version management complexities and promotes a safer development and testing process by creating a dedicated development environment.

upvoted 4 times

🗨️ **tttamttam** 1 year ago

Selected Answer: BD

D ==> change the lambda function.

B ==> update the API gateway to use the updated lambda function & deploy it into another(new) stage. so that developers can use the newly deployed API endpoint.

upvoted 4 times

🗨️ **csG13** 1 year, 1 month ago

Selected Answer: BD

It is B & D.

Clearly E isn't operationally efficient. So we got to choose from A & B one, and C & D the second.

Between A & B, we gotta pick B since in the question it clearly states that we don't want to touch the existing solution.

Regarding C & D, seems like D is more thorough and also pointing to \$LATEST is not sufficiently explicit when you troubleshoot.

upvoted 4 times

🗨️ **zodraz** 1 year, 2 months ago

Selected Answer: BD

<https://www.examttopics.com/discussions/amazon/view/89549-exam-aws-certified-developer-associate-topic-1-question-334/>

upvoted 2 times

A developer is implementing an AWS Cloud Development Kit (AWS CDK) serverless application. The developer will provision several AWS Lambda functions and Amazon API Gateway APIs during AWS CloudFormation stack creation. The developer's workstation has the AWS Serverless Application Model (AWS SAM) and the AWS CDK installed locally.

How can the developer test a specific Lambda function locally?

- A. Run the sam package and sam deploy commands. Create a Lambda test event from the AWS Management Console. Test the Lambda function.
- B. Run the cdk synth and cdk deploy commands. Create a Lambda test event from the AWS Management Console. Test the Lambda function.
- C. Run the cdk synth and sam local invoke commands with the function construct identifier and the path to the synthesized CloudFormation template.
- D. Run the cdk synth and sam local start-lambda commands with the function construct identifier and the path to the synthesized CloudFormation template.

Correct Answer: C

Community vote distribution

C (100%)

MrTee **Highly Voted** 1 year, 9 months ago

Selected Answer: C

The developer can test a specific Lambda function locally by running the cdk synth command to synthesize the AWS CDK application into an AWS CloudFormation template. Then, the developer can use the sam local invoke command with the function construct identifier and the path to the synthesized CloudFormation template to test the Lambda function locally (option C).

upvoted 13 times

Dushank **Highly Voted** 1 year, 4 months ago

Selected Answer: C

To test a specific Lambda function locally when using the AWS Cloud Development Kit (AWS CDK), the developer can use the AWS Serverless Application Model (AWS SAM) CLI's local testing capabilities in conjunction with the CDK. The typical process would be:

Run cdk synth to synthesize the AWS CDK app into a CloudFormation template.

Use sam local invoke to run the specific Lambda function locally, providing the function's logical identifier and the path to the synthesized CloudFormation template as arguments.

upvoted 5 times

sumanshu **Most Recent** 1 month ago

Selected Answer: C

A) Eliminated - The sam package and sam deploy commands are used to upload the Lambda function code and deploy it to AWS.

B) Eliminated - The cdk deploy command deploys the stack to AWS.

C) Correct - The sam local invoke command uses this template to run the Lambda function locally, simulating the AWS Lambda execution environment.

D) Eliminated

upvoted 1 times

nbxyzd 3 months ago

Selected Answer: C

Note that the sam local start-lambda subcommand starts a local endpoint to emulate AWS Lambda.

upvoted 2 times

Anandesh 6 months, 1 week ago

Selected Answer: C

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/serverless-cdk-testing.html>

upvoted 1 times

🗨️ 👤 **65703c1** 8 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗨️ 👤 **KillThemWithKindness** 11 months, 2 weeks ago

Selected Answer: C

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/serverless-sam-cli-using-invoke.html>

sam local invoke: Invoke Lambda locally

sam local start-lambda: Integrating with automated-tests

upvoted 1 times

🗨️ 👤 **NaghamAbdellatif** 1 year, 2 months ago

Selected Answer: C

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/serverless-cdk-testing.html>

C

upvoted 1 times

🗨️ 👤 **fossil123** 1 year, 5 months ago

Selected Answer: C

Use the AWS SAM CLI sam local invoke subcommand to initiate a one-time invocation of an AWS Lambda function locally.

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/using-sam-cli-local-invoke.html>

upvoted 2 times

🗨️ 👤 **JamalDaBoss** 1 year, 6 months ago

Selected Answer: C

Answer is clearly C. If you say it's not C, you are wrong.

upvoted 2 times

🗨️ 👤 **zodraz** 1 year, 8 months ago

Selected Answer: C

sam local invoke StackLogicalId/FunctionLogicalId

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/using-sam-cli-local-invoke.html>

upvoted 4 times

A company's new mobile app uses Amazon API Gateway. As the development team completes a new release of its APIs, a developer must safely and transparently roll out the API change.

What is the SIMPLEST solution for the developer to use for rolling out the new API version to a limited number of users through API Gateway?

- A. Create a new API in API Gateway. Direct a portion of the traffic to the new API using an Amazon Route 53 weighted routing policy.
- B. Validate the new API version and promote it to production during the window of lowest expected utilization.
- C. Implement an Amazon CloudWatch alarm to trigger a rollback if the observed HTTP 500 status code rate exceeds a predetermined threshold.
- D. Use the canary release deployment option in API Gateway. Direct a percentage of the API traffic using the canarySettings setting.

Correct Answer: D

Community vote distribution

D (100%)

  **zodraz** Highly Voted 1 year, 2 months ago



Selected Answer: D

<https://www.examttopics.com/discussions/amazon/view/51596-exam-aws-certified-developer-associate-topic-1-question-355/>
upvoted 7 times

  **sumanshu** Most Recent 1 month ago



Selected Answer: D

- A) Eliminated - creating a new API in API Gateway requires managing two separate APIs, which adds unnecessary complexity.
- B) Eliminated - This is a direct cutover approach, not a gradual rollout to a limited number of users. If issues arise, all users are impacted, which is not safe or transparent.
- C) Eliminated - It is not relevant to the question's requirement
- D) Correct - It allows gradual exposure of the new API version to a subset of users, minimizing risk.
upvoted 1 times

  **65703c1** 2 months, 1 week ago

Selected Answer: D

D is the correct answer.
upvoted 1 times

  **Dushank** 10 months, 3 weeks ago

Selected Answer: D

Canary deployments allow you to divert a percentage of your API traffic to a new API version, enabling you to test how the new version will perform under real-world conditions without fully replacing the previous version. This is especially useful for reducing the risk associated with deploying new versions.
upvoted 4 times

A company caches session information for a web application in an Amazon DynamoDB table. The company wants an automated way to delete old items from the table.

What is the simplest way to do this?

- A. Write a script that deletes old records; schedule the script as a cron job on an Amazon EC2 instance.
- B. Add an attribute with the expiration time; enable the Time To Live feature based on that attribute.
- C. Each day, create a new table to hold session data; delete the previous day's table.
- D. Add an attribute with the expiration time; name the attribute ItemExpiration.

Correct Answer: B

Community vote distribution

B (100%)

 **Dushank** Highly Voted 10 months, 3 weeks ago

Selected Answer: B

The simplest way to automatically delete old items from an Amazon DynamoDB table is to use DynamoDB's Time to Live (TTL) feature. This feature allows you to define an attribute that stores the expiration time for each item. Once the specified time has passed, DynamoDB automatically deletes the expired items, freeing up storage and reducing costs without the need for custom scripts or manual intervention.


upvoted 6 times

 **sumanshu** Most Recent 1 month ago

Selected Answer: B

- A) Eliminated - no automation
- C) Eliminated - Constantly creating and deleting tables is not practical and can lead to management and scaling issues.
- D) Eliminated - Without enabling TTL, merely adding the attribute does not automate item deletion.

upvoted 1 times

 **65703c1** 2 months, 1 week ago

Selected Answer: B


B is the correct answer.

upvoted 1 times

 **catcatpunch** 1 year, 2 months ago

https://docs.aws.amazon.com/ko_kr/amazondynamodb/latest/developerguide/TTL.html

upvoted 2 times

 **zodraz** 1 year, 2 months ago

Selected Answer: B

<https://www.examttopics.com/discussions/amazon/view/7225-exam-aws-certified-developer-associate-topic-1-question-107/>

upvoted 4 times

A company is using an Amazon API Gateway REST API endpoint as a webhook to publish events from an on-premises source control management (SCM) system to Amazon EventBridge. The company has configured an EventBridge rule to listen for the events and to control application deployment in a central AWS account. The company needs to receive the same events across multiple receiver AWS accounts.

How can a developer meet these requirements without changing the configuration of the SCM system?

- A. Deploy the API Gateway REST API to all the required AWS accounts. Use the same custom domain name for all the gateway endpoints so that a single SCM webhook can be used for all events from all accounts.
- B. Deploy the API Gateway REST API to all the receiver AWS accounts. Create as many SCM webhooks as the number of AWS accounts.
- C. Grant permission to the central AWS account for EventBridge to access the receiver AWS accounts. Add an EventBridge event bus on the receiver AWS accounts as the targets to the existing EventBridge rule.
- D. Convert the API Gateway type from REST API to HTTP API.

Correct Answer: C

Community vote distribution

C (100%)

 **csG13** Highly Voted 1 year, 1 month ago

Selected Answer: C

It's C - eventbridge event buses in one (target) account can be a target of another event rule in a source account.

For reference, watch the video in the following link:

<https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-cross-account.html>


upvoted 19 times

 **sumanshu** Most Recent 1 month ago

Selected Answer: C

This approach uses EventBridge's native cross-account event delivery capabilities, ensuring events are sent to the receiver accounts without modifying the SCM system configuration.

upvoted 1 times

 **65703c1** 2 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times


A company moved some of its secure files to a private Amazon S3 bucket that has no public access. The company wants to develop a serverless application that gives its employees the ability to log in and securely share the files with other users.

Which AWS feature should the company use to share and access the files securely?

- A. Amazon Cognito user pool
- B. S3 presigned URLs
- C. S3 bucket policy
- D. Amazon Cognito identity pool

Correct Answer: B

Community vote distribution



🗳️ **Dushank** Highly Voted 1 year, 4 months ago

Selected Answer: B

Employees log into the serverless application using an Amazon Cognito User Pool.

Once logged in, the application's back-end logic (possibly a Lambda function) generates an S3 pre-signed URL for the requested file.

The pre-signed URL is then given to the authenticated user, allowing them secure, time-limited access to that specific S3 object.

So, while both Amazon Cognito User Pool and S3 Pre-signed URLs would be used in the solution, S3 Pre-signed URLs (Option B) are the specific feature that allows for the secure, temporary sharing of S3 files.

Therefore, Option B would be the best answer to the question of how to "share and access the files securely."

upvoted 25 times

🗳️ **loctong** Highly Voted 1 year, 8 months ago

Selected Answer: A

the key words are ability to log in and securely share the files. It is A

upvoted 20 times

🗳️ **jipark** 1 year, 6 months ago

I agree 'log in' would go user pool.

upvoted 2 times

🗳️ **rimaSamir** 11 months, 2 weeks ago

But we need to answer a question not task condition

upvoted 3 times

🗳️ **sumanshu** Most Recent 1 month ago

Selected Answer: B

A) Eliminated - While Cognito is useful for managing user authentication, it does not directly provide the capability to securely share files from an S3 bucket.

upvoted 1 times

🗳️ **wh1t4k3r** 5 months, 2 weeks ago

Selected Answer: B

I had my doubt on this, but once more they are evaluating if you are reading with attention and not if you have knowledge hehe tricky question, but the punch line question is What feature would be used to share the files securely, ignoring the login part.

upvoted 4 times

🗳️ **kyowo** 6 months, 4 weeks ago

Selected Answer: B

I'll go with B. The question is what the company should use to share and access the files securely. We can ignore the task condition

upvoted 3 times

🗳️ **65703c1** 8 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

SerialiDr 11 months, 1 week ago

Selected Answer: B

This option allows secure, temporary access to specific objects in an S3 bucket. By generating presigned URLs, the serverless application can grant users time-limited access to download or upload files without altering the permissions of the S3 bucket or the objects. This method ensures secure access management and is suitable for sharing private files among authenticated users.

upvoted 3 times

SD_CS 12 months ago

Selected Answer: A

in order to log in you need to use cognito user pools

upvoted 3 times

rimaSamir 1 year ago

Actually, the question is about "what feature will be used by the new serverless application to share and access the files securely". Ability to log in is about "Amazon Cognito user pool". Imagine "Lambda function" and "API Gateway" are created as a serverless app to provide some API. When you call API endpoint, it will login to "Amazon Cognito user pool" and then share files using SDK. How it will share is the next question.

My answer is A

upvoted 4 times

Ashwinvdm22 1 year ago

Selected Answer: B

The answer must be B. So although in the question it says "gives its employees the ability to log in" (which is hinting towards Cognito User Pools) the question is actually asking: "Which AWS feature should the company use to share and access the files securely?"

The question is actually about how to share and access the files securely. Hence it must be the S3 pre-signed URL option. To read up more on S3 pre-signed URLs check here: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/ShareObjectPreSignedURL.html>

upvoted 4 times

peekingpicker 1 year ago

Selected Answer: B

Which AWS feature should the company use to share and access the files securely?

So, It's B. S3 Pre-signed URL can used to share S3 object to other people securely.

upvoted 3 times

gqs3119 1 year, 1 month ago

It's not A, Cognito user pool is not needed, only employees need ability to log in, they can be provided with IAM accounts.

upvoted 1 times

a_win 1 year, 1 month ago

Selected Answer: D

An Amazon Cognito identity pool provides temporary AWS credentials for users who authenticate via Amazon Cognito. This allows your application users (employees, in this case) to securely authenticate and gain access to AWS services like S3 based on their assigned roles and permissions.

Through Amazon Cognito, you can manage user identities, control user access to resources, and provide temporary, limited-privilege credentials to access the S3 bucket securely.

upvoted 3 times

KarBiswa 1 year, 1 month ago

Selected Answer: B

I will go with B because its purely asking about sharing and no mention about external logins so we should go by default AWS feature which provides this feature,

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/ShareObjectPreSignedURL.html>

upvoted 3 times

tqiu654 1 year, 2 months ago

Selected Answer: B

ChatGPT: B

upvoted 3 times

didorins 1 year, 3 months ago

Login of external to AWS users, we can use Cognito. Identity Pool is specifically for DynamoDB and S3.

Use an identity pool when you need to:

Give your users access to AWS resources, such as an Amazon Simple Storage Service (Amazon S3) bucket or an Amazon DynamoDB table.

<https://repost.aws/knowledge-center/cognito-user-pools-identity-pools>

upvoted 1 times

  **Rameez1** 1 year, 3 months ago

Selected Answer: B

Actual ask is in the final line "Which AWS feature should the company use to share and access the files securely?" -> S3 Pre-signed URL provides the most secure feature.

upvoted 4 times

  **[Removed]** 1 year, 1 month ago

I agree... B is the only option that is specific to sharing of files. Identity/User pools are for authentication (log in to the S3 bucket).

upvoted 3 times

A company needs to develop a proof of concept for a web service application. The application will show the weather forecast for one of the company's office locations. The application will provide a REST endpoint that clients can call. Where possible, the application should use caching features provided by AWS to limit the number of requests to the backend service. The application backend will receive a small amount of traffic only during testing.

Which approach should the developer take to provide the REST endpoint MOST cost-effectively?

- A. Create a container image. Deploy the container image by using Amazon Elastic Kubernetes Service (Amazon EKS). Expose the functionality by using Amazon API Gateway.
- B. Create an AWS Lambda function by using the AWS Serverless Application Model (AWS SAM). Expose the Lambda functionality by using Amazon API Gateway.
- C. Create a container image. Deploy the container image by using Amazon Elastic Container Service (Amazon ECS). Expose the functionality by using Amazon API Gateway.
- D. Create a microservices application. Deploy the application to AWS Elastic Beanstalk. Expose the AWS Lambda functionality by using an Application Load Balancer.

Correct Answer: B

Community vote distribution

B (100%)

 **loctong** Highly Voted 1 year, 2 months ago

Selected Answer: B

AWS Lambda function absolutely ability to do the requirements.

upvoted 9 times

 **JamalDaBoss** 12 months ago

Yes, Lambda bery certain great.

upvoted 3 times

 **sumanshu** Most Recent 1 month ago


Selected Answer: B

A) Eliminated - Amazon EKS is designed for large-scale, containerized applications that require Kubernetes orchestration. For a simple proof of concept with small traffic, it is overkill and introduces unnecessary complexity and cost.

C) Eliminated - ECS introduces additional operational overhead compared to a serverless solution like Lambda.

D) Elastic Beanstalk is not designed for a lightweight serverless application; it requires provisioning and managing compute resources, which increases cost and complexity.

upvoted 1 times

 **65703c1** 2 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

 **SerialiDr** 5 months ago

Selected Answer: B

This solution is cost-effective because AWS Lambda charges are based on the number of requests and the duration of code execution, making it ideal for applications with low to moderate traffic. Amazon API Gateway can efficiently manage the REST endpoint and offers built-in caching capabilities to reduce the number of requests to the backend Lambda function, further optimizing costs. This setup also leverages the serverless model, reducing the operational overhead and cost associated with provisioning and managing servers.



upvoted 4 times

 **a5fc516** 5 months, 4 weeks ago

Selected Answer: B

yes B is correct

upvoted 1 times

  **hmdev** 11 months, 1 week ago

Selected Answer: B

B is the cost-effective one.

upvoted 3 times

An e-commerce web application that shares session state on-premises is being migrated to AWS. The application must be fault tolerant, natively highly scalable, and any service interruption should not affect the user experience.

What is the best option to store the session state?

- A. Store the session state in Amazon ElastiCache.
- B. Store the session state in Amazon CloudFront.
- C. Store the session state in Amazon S3.
- D. Enable session stickiness using elastic load balancers.

Correct Answer: A

Community vote distribution

A (100%)

🗨️ **sumanshu** 1 month ago

Selected Answer: A

B) Eliminated - Amazon CloudFront is a content delivery network (CDN) designed for distributing static and dynamic content globally.

C) Eliminated - it is not designed for low-latency use cases like session management.

D) Eliminated - it is not highly scalable, as it ties user sessions to specific servers, which can lead to uneven load distribution and scaling issues.
upvoted 1 times

🗨️ **Saudis** 4 months, 2 weeks ago

Selected Answer: A

A is Correct, the store the session state is a key word
upvoted 1 times

🗨️ **65703c1** 8 months, 1 week ago

Selected Answer: A

A is the correct answer.
upvoted 2 times

🗨️ **SerialiDr** 11 months, 1 week ago

Selected Answer: A

Amazon ElastiCache is a high-performance, in-memory data store that provides sub-millisecond latency to applications. It supports data structures such as strings, hashes, lists, sets, and sorted sets, making it suitable for storing session state data. ElastiCache offers both Redis and Memcached engines, with Redis providing more advanced data structures and features such as persistence, replication, and transaction support. This solution is fault-tolerant and highly scalable, ensuring that any service interruption does not affect the user experience.
upvoted 4 times

🗨️ **Phongsanth** 1 year, 7 months ago

Selected Answer: A

I vote A

<https://aws.amazon.com/blogs/developer/elasticache-as-an-asp-net-session-store/>
upvoted 4 times

🗨️ **loctong** 1 year, 8 months ago

Selected Answer: A

the answer came from the discussion at <https://www.examtopycs.com/discussions/amazon/view/8789-exam-aws-certified-developer-associate-topic-1-question-176/>
upvoted 3 times

🗨️ **zodraz** 1 year, 8 months ago

Selected Answer: A

<https://www.examttopics.com/discussions/amazon/view/8789-exam-aws-certified-developer-associate-topic-1-question-176/>
upvoted 4 times

A developer is building an application that uses Amazon DynamoDB. The developer wants to retrieve multiple specific items from the database with a single API call.

Which DynamoDB API call will meet these requirements with the MINIMUM impact on the database?

- A. BatchGetItem
- B. GetItem
- C. Scan
- D. Query

Correct Answer: A

Community vote distribution

A (100%)

MrTee **Highly Voted** 1 year, 9 months ago

Selected Answer: A

A is the correct answer with the minimum impact on the database.

upvoted 12 times

dan80 **Highly Voted** 1 year, 9 months ago

Selected Answer: A

<https://beabetterdev.com/2022/10/12/dynamodb-getitem-vs-query-when-to-use-what/#:~:text=If%20you'd%20like%20to%20retrieve%20multiple%20items%20at%20once,retrieve%20multiple%20items%20at%20once.>

upvoted 9 times

jipark 1 year, 6 months ago

tons of thanks.

Looking for just a single item on the main table index? Use GetItem

Looking for just a single item on a GSI? Use Query.

Looking for multiple items with different partition key and sort key combinations at once? Use BatchGetItem

Looking for multiple items that share the same partition key? Use Query

upvoted 14 times

sumanshu **Most Recent** 1 month ago

Selected Answer: A

B) Eliminated - The GetItem API retrieves only one specific item at a time by its primary key.

C) Eliminated - The Scan API retrieves all items in a table or index by sequentially scanning every item.

D) Eliminated - The Query API retrieves all items that match a specific partition key and optionally filters them by sort key or other criteria.

upvoted 1 times

Saudis 4 months, 2 weeks ago

A key word is retrieve multiple specific items

upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

Moralles 9 months, 3 weeks ago

Selected Answer: A

• GetItem: recupera um único item de uma tabela.

• BatchGetItem: retorna até 100 itens de uma ou mais tabelas limitados a 16MB por chamada.

upvoted 1 times

  **prathameshpathak** 1 year ago

Selected Answer: A

.....

upvoted 1 times

  **marolisa** 1 year, 4 months ago

D.

"Query" allows you to use filter - multiple specific items and is less expensive than the Scan operation.



upvoted 1 times

  **Baba_Eni** 1 year, 7 months ago

Selected Answer: A

https://docs.aws.amazon.com/amazondynamodb/latest/APIReference/API_BatchGetItem.html

upvoted 1 times

  **imvb88** 1 year, 8 months ago

Selected Answer: A

Need specific Item -> cannot be Scan or Query since they are for retrieving items that match conditions.

We need multiple item then A is the option left.

upvoted 2 times

A developer has written an application that runs on Amazon EC2 instances. The developer is adding functionality for the application to write objects to an Amazon S3 bucket.

Which policy must the developer modify to allow the instances to write these objects?

- A. The IAM policy that is attached to the EC2 instance profile role
- B. The session policy that is applied to the EC2 instance role session
- C. The AWS Key Management Service (AWS KMS) key policy that is attached to the EC2 instance profile role
- D. The Amazon VPC endpoint policy

Correct Answer: A

Community vote distribution

A (100%)

🗨️ **Prem28** Highly Voted 1 year, 8 months ago

Selected Answer: A

a is correct

upvoted 6 times

🗨️ **Ja13** Highly Voted 1 year, 7 months ago

Selected Answer: A

A: <https://repost.aws/knowledge-center/ec2-instance-access-s3-bucket>

upvoted 5 times

🗨️ **sumanshu** Most Recent 1 month ago

Selected Answer: A

EC2 instances assume an IAM role through an instance profile.

upvoted 1 times

🗨️ **albert_kuo** 4 months, 1 week ago

Selected Answer: A

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
        "s3:PutObjectAcl"
      ],
      "Resource": "arn:aws:s3:::your-bucket-name/*"
    }
  ]
}
```

upvoted 1 times

🗨️ **65703c1** 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ **mgonblan** 1 year, 8 months ago

B: I Think B is better, because we need to use it on the instance session

upvoted 1 times

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. VPN logs
- B. BGP logs
- C. VPC Flow Logs
- D. AWS CloudTrail logs

Correct Answer: C

Community vote distribution

C (100%)

 **Dushank** Highly Voted 1 year, 4 months ago

Selected Answer: C

VPC Flow Logs capture information about the IP traffic going to and from network interfaces in a VPC. This includes traffic that traverses a VPN connection. VPC Flow Logs can be used to monitor and troubleshoot connectivity issues, including verifying whether traffic is reaching a particular subnet within the VPC.

upvoted 8 times

 **sumanshu** Most Recent 1 month ago

Selected Answer: C

VPC Flow Logs capture detailed information about the traffic flowing to and from network interfaces in a VPC.

D) Eliminated - AWS CloudTrail logs record API-level activity in an AWS account, they do not provide network-level traffic details for verifying traffic reaching specific subnets.

upvoted 1 times

 **albert_kuo** 4 months, 1 week ago


Selected Answer: C

version account-id interface-id srcaddr dstaddr srcport dstport protocol packets bytes start end action log-status

2 123456789012 eni-abc12345 192.168.1.10 10.0.0.20 443 80 6 10 2048 1620050730 1620050790 ACCEPT OK

2 123456789012 eni-abc12345 192.168.1.10 10.0.0.30 22 443 6 5 1024 1620050730 1620050790 REJECT OK

upvoted 1 times

 **65703c1** 8 months, 1 week ago

Selected Answer: C

C is the correct answer.


upvoted 1 times

 **KarBiswa** 1 year, 1 month ago

Selected Answer: C

<https://docs.aws.amazon.com/vpc/latest/userguide/flow-logs.html>


upvoted 1 times

 **Prem28** 1 year, 8 months ago

Selected Answer: C

<https://www.examttopics.com/discussions/amazon/view/28802-exam-aws-certified-developer-associate-topic-1-question-219/>

upvoted 3 times

 **zodraz** 1 year, 8 months ago

Selected Answer: C

<https://www.examttopics.com/discussions/amazon/view/28802-exam-aws-certified-developer-associate-topic-1-question-219/>

upvoted 3 times

A developer is creating a service that uses an Amazon S3 bucket for image uploads. The service will use an AWS Lambda function to create a thumbnail of each image. Each time an image is uploaded, the service needs to send an email notification and create the thumbnail. The developer needs to configure the image processing and email notifications setup.

Which solution will meet these requirements?

- A. Create an Amazon Simple Notification Service (Amazon SNS) topic. Configure S3 event notifications with a destination of the SNS topic. Subscribe the Lambda function to the SNS topic. Create an email notification subscription to the SNS topic.
- B. Create an Amazon Simple Notification Service (Amazon SNS) topic. Configure S3 event notifications with a destination of the SNS topic. Subscribe the Lambda function to the SNS topic. Create an Amazon Simple Queue Service (Amazon SQS) queue. Subscribe the SQS queue to the SNS topic. Create an email notification subscription to the SQS queue.
- C. Create an Amazon Simple Queue Service (Amazon SQS) queue. Configure S3 event notifications with a destination of the SQS queue. Subscribe the Lambda function to the SQS queue. Create an email notification subscription to the SQS queue.
- D. Create an Amazon Simple Queue Service (Amazon SQS) queue. Send S3 event notifications to Amazon EventBridge. Create an EventBridge rule that runs the Lambda function when images are uploaded to the S3 bucket. Create an EventBridge rule that sends notifications to the SQS queue. Create an email notification subscription to the SQS queue.

Correct Answer: A

Community vote distribution

A (100%)

MrTee Highly Voted 1 year, 9 months ago

Selected Answer: A

This solution will allow the developer to receive notifications for each image uploaded to the S3 bucket, and also create a thumbnail using the Lambda function. The SNS topic will serve as a trigger for both the Lambda function and the email notification subscription. When an image is uploaded, S3 will send a notification to the SNS topic, which will trigger the Lambda function to create the thumbnail and also send an email notification to the specified email address.

upvoted 16 times

payireb682 1 year, 1 month ago

Thanks. As mentioned Multiple subscription can be added for SNS

upvoted 5 times

jipark 1 year, 6 months ago

greate !! send email do not need SQS.

upvoted 2 times

sumanshu Most Recent 1 month ago

Selected Answer: A

B) Eliminated - The SQS queue is not required because SNS can directly notify both the Lambda function and the email service

C) Eliminated - SQS cannot send email notifications. Also, The Lambda function would need to poll the SQS queue, adding unnecessary complexity.

D) Eliminated - Using EventBridge to forward S3 event notifications introduces unnecessary complexity.

upvoted 1 times

Tluszczyk 6 months, 2 weeks ago

None of these is really an optimal solution to the problem, which is a little annoying really

upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

  **SerialIDr** 1 year ago

Selected Answer: A

SNS can be used to fan out notifications. When an image is uploaded to the S3 bucket, an event notification is sent to the SNS topic. The Lambda function is subscribed to this topic to create a thumbnail, and an email subscription can also be configured on the same SNS topic to send email notifications. This approach meets all requirements with minimal components.

upvoted 2 times

A developer has designed an application to store incoming data as JSON files in Amazon S3 objects. Custom business logic in an AWS Lambda function then transforms the objects, and the Lambda function loads the data into an Amazon DynamoDB table. Recently, the workload has experienced sudden and significant changes in traffic. The flow of data to the DynamoDB table is becoming throttled.


The developer needs to implement a solution to eliminate the throttling and load the data into the DynamoDB table more consistently.

Which solution will meet these requirements?

- A. Refactor the Lambda function into two functions. Configure one function to transform the data and one function to load the data into the DynamoDB table. Create an Amazon Simple Queue Service (Amazon SQS) queue in between the functions to hold the items as messages and to invoke the second function.
- B. Turn on auto scaling for the DynamoDB table. Use Amazon CloudWatch to monitor the table's read and write capacity metrics and to track consumed capacity.
- C. Create an alias for the Lambda function. Configure provisioned concurrency for the application to use.
- D. Refactor the Lambda function into two functions. Configure one function to store the data in the DynamoDB table. Configure the second function to process the data and update the items after the data is stored in DynamoDB. Create a DynamoDB stream to invoke the second function after the data is stored.

Correct Answer: A

Community vote distribution



ihebchorfi Highly Voted 1 year, 9 months ago

Selected Answer: A

A. Refactor the Lambda function into two functions. Configure one function to transform the data and one function to load the data into the DynamoDB table. Create an Amazon Simple Queue Service (Amazon SQS) queue in between the functions to hold the items as messages and to invoke the second function.

By breaking the Lambda function into two separate functions and using an SQS queue to hold the transformed data as messages, you can decouple the data transformation and loading processes. This allows for more controlled loading of data into the DynamoDB table and helps eliminate throttling issues.

upvoted 26 times

MrTee Highly Voted 1 year, 9 months ago

Selected Answer: D

This solution will allow the developer to store the incoming data into the DynamoDB table more consistently without being throttled. By splitting the Lambda function into two functions, the first function can store the data into the DynamoDB table and exit quickly, avoiding any throttling issues. The second function can then process the data and update the items after the data is stored in DynamoDB using a DynamoDB stream to invoke the second function.

Option A is also a good option but not the best solution because it introduces additional complexity and cost by using an Amazon SQS queue.

upvoted 9 times

nbxyzd 3 months ago

Read carefully: The flow of data to the DynamoDB table is becoming throttled.

So the bottleneck is the DynamoDB, not the lambda function transforming the data. Option D doesn't help because the first function storing data into the DynamoDB will still hit the throttling issue.

upvoted 2 times

Ashwinvdm22 1 year ago


The problem I have with option D is that it is adding more load on the DynamoDB table. What is the need to insert the item and then update the item later. This is performing two operation on every item just to get it into the correct state. I would go with option A since it is not performing two operations on the DB and hence reducing the load which will help with throttling.

upvoted 1 times

robotgeek 1 year, 8 months ago

Sorry but when you say "the first function can store the data into the DynamoDB table and exit quickly, avoiding any throttling issues" I dont understand your point

upvoted 5 times

  **[Removed]** 1 year, 1 month ago

I disagree... the order of the function with this option makes NO sense. I go with A

upvoted 1 times

  **mooncake1** Most Recent 6 days, 22 hours ago

Selected Answer: B

Adjusting the WCU is enough. no need to implement such complex solutions

upvoted 1 times

  **sumanshu** 1 month ago

Selected Answer: A

The SQS queue acts as a buffer, which smooths out sudden traffic spikes by queuing up data. The second Lambda function processes data from the queue at a steady rate, reducing the likelihood of throttling in DynamoDB.

B) Eliminated - Auto scaling takes time to adjust, during which throttling may still occur.

C) Eliminated - Provisioned concurrency ensures that the Lambda function can handle a predictable number of concurrent requests, but it does not solve the throttling issue at the DynamoDB layer



upvoted 1 times

  **nbxyzd** 3 months ago

Selected Answer: A

It's hands-down A.

upvoted 1 times



  **Tluszczuk** 6 months, 2 weeks ago

Selected Answer: B

"A" would be optimal, but without backoff algorithm the lambda division and SQS won't affect the throttling. However Dynamo can autoscale


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

upvoted 1 times

  **queekao** 6 months, 1 week ago

Yes, but it asks to "load the data into the DynamoDB table more consistently." Therefore, option A will prevent unintentional data load into DynamoDB, it's the best option.

upvoted 1 times

  **65703c1** 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

  **nder** 10 months, 3 weeks ago

Selected Answer: B

we are trying to stop throttling...

upvoted 1 times

  **SerialiDr** 11 months, 1 week ago

Selected Answer: A

This solution addresses the need to eliminate throttling and ensure consistent data loading into the Amazon DynamoDB table by separating the transformation and loading processes into two different functions. Using an Amazon SQS queue to hold items as messages between the two functions helps manage the flow of data and prevents overloading the DynamoDB table, thereby eliminating throttling issues.

upvoted 2 times

  **Brisun** 12 months ago

Selected Answer: A

A is correct as it requires to write to DynamoDB "more consistently". Option B can solve the problem too but the writing won't be consistent as the traffic will go up and down instantly.

In reality, I will probably do Option B only.

upvoted 3 times

SD_CS 12 months ago

Selected Answer: B

I do not feel refactoring the data transformation and loading would help here as I do not think the number of concurrent calls to the DB would decrease because of this. Autoscaling DynamoDB would seem a more potent option to me.

upvoted 4 times

peekingpicker 1 year ago

Selected Answer: B

Why not B ?

DynamoDB can autoscale the RCU and WCU

upvoted 4 times

SerialiDr 1 year ago

Selected Answer: A

A. Refactor the Lambda function into two functions, using an Amazon SQS queue to manage the data flow, and/or

B. Turn on auto scaling for the DynamoDB table to automatically adjust its write capacity based on traffic patterns.

Both A and B address the core issue of managing write throughput to the DynamoDB table to prevent throttling. Option A provides a way to smooth out data flow and manage write requests more effectively, while option B allows the table to scale its capacity automatically in response to changing traffic, although with potential limitations in response speed to sudden traffic spikes. Combining these approaches could provide an even more robust solution.

upvoted 3 times

KarBiswa 1 year, 1 month ago

Selected Answer: A

Off course A & D are options but here after inserting the data further we cannot modify because one extra writing cost will incur rather using queue lambda can poll the transformed data

upvoted 3 times

Nagasoacle 1 year, 3 months ago

Selected Answer: A

Answer : A

SQS can be configured to invoke Lambda.

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-configure-lambda-function-trigger.html>

upvoted 4 times

dexdinh91 1 year, 3 months ago

Selected Answer: B

I think B

upvoted 1 times

jingle4944 1 year, 3 months ago

Lambda functions can be triggered by SQS: <https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-configure-lambda-function-trigger.html>

upvoted 2 times

nbxyzd 3 months ago

Correct, and most importantly, it's triggered *synchronously* so that there won't be throttling issue. Quote:

You can use an AWS Lambda function to process messages in an Amazon SQS queue. Lambda polls the queue and invokes your Lambda function synchronously with an event that contains queue messages.

upvoted 1 times

A developer is creating an AWS Lambda function in VPC mode. An Amazon S3 event will invoke the Lambda function when an object is uploaded into an S3 bucket. The Lambda function will process the object and produce some analytic results that will be recorded into a file. Each processed object will also generate a log entry that will be recorded into a file.

Other Lambda functions, AWS services, and on-premises resources must have access to the result files and log file. Each log entry must also be appended to the same shared log file. The developer needs a solution that can share files and append results into an existing file.

Which solution should the developer use to meet these requirements?

- A. Create an Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system in Lambda. Store the result files and log file in the mount point. Append the log entries to the log file.
- B. Create an Amazon Elastic Block Store (Amazon EBS) Multi-Attach enabled volume. Attach the EBS volume to all Lambda functions. Update the Lambda function code to download the log file, append the log entries, and upload the modified log file to Amazon EBS.
- C. Create a reference to the /tmp local directory. Store the result files and log file by using the directory reference. Append the log entry to the log file.
- D. Create a reference to the /opt storage directory. Store the result files and log file by using the directory reference. Append the log entry to the log file.

Correct Answer: A

Community vote distribution

A (100%)

 **Dushank** Highly Voted 1 year, 4 months ago

Selected Answer: A

The requirement is to have a shared file system that allows for appending to files and can be accessed by multiple Lambda functions, AWS services, and on-premises resources. Amazon Elastic File System (Amazon EFS) is a good fit for these requirements. EFS provides a scalable and elastic NFS file system which can be mounted to multiple EC2 instances and Lambda functions at the same time, making it easier for these resources to share files. You can also append to existing files on an EFS file system, which meets the requirement for a shared log file that can have new entries appended to it.

upvoted 9 times

 **sumanshu** Most Recent 1 month ago

Selected Answer: A

B) Eliminated - Amazon Elastic Block Store (EBS) volumes are designed for single-instance use or Multi-Attach for specific EC2 instances, not for Lambda functions or serverless environments.

C) Eliminated - The /tmp directory in Lambda is ephemeral storage that exists only for the duration of the Lambda invocation. It is not shared between functions or persistent across invocations

D) Eliminated - The /opt directory in Lambda is used for deploying custom runtimes or external libraries. It is read-only during execution, so it cannot be used to store or update files.


upvoted 1 times

 **Saudis** 4 months, 3 weeks ago

Selected Answer: A

Elastic File System file shareing

upvoted 1 times

 **65703c1** 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

 **mgonblan** 1 year, 8 months ago



A) There are several references for this:

<https://docs.aws.amazon.com/lambda/latest/operatorguide/networking-vpc.html> and

this blog entry:

<https://aws.amazon.com/es/blogs/compute/choosing-between-aws-lambda-data-storage-options-in-web-apps/>



upvoted 1 times

  **delak** 1 year, 8 months ago

Selected Answer: A

shared files == EFS

upvoted 3 times

  **loctong** 1 year, 8 months ago

Selected Answer: A

EFS is true

upvoted 2 times

A company has an AWS Lambda function that processes incoming requests from an Amazon API Gateway API. The API calls the Lambda function by using a Lambda alias. A developer updated the Lambda function code to handle more details related to the incoming requests. The developer wants to deploy the new Lambda function for more testing by other developers with no impact to customers that use the API.

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

- A. Create a new version of the Lambda function. Create a new stage on API Gateway with integration to the new Lambda version. Use the new API Gateway stage to test the Lambda function.
- B. Update the existing Lambda alias used by API Gateway to a weighted alias. Add the new Lambda version as an additional Lambda function with a weight of 10%. Use the existing API Gateway stage for testing.
- C. Create a new version of the Lambda function. Create and deploy a second Lambda function to filter incoming requests from API Gateway. If the filtering Lambda function detects a test request, the filtering Lambda function will invoke the new Lambda version of the code. For other requests, the filtering Lambda function will invoke the old Lambda version. Update the API Gateway API to use the filtering Lambda function.
- D. Create a new version of the Lambda function. Create a new API Gateway API for testing purposes. Update the integration of the new API with the new Lambda version. Use the new API for testing.

Correct Answer: A

Community vote distribution



🗨️ **sumanshu** 1 month ago

Selected Answer: A

Existing customers continue to use the current stage (e.g., prod), which is integrated with the old Lambda version. The new stage (e.g., test) can point to the new Lambda version without affecting the production stage.

B) Eliminated - routing even 10% of traffic to the new Lambda version could cause issues for customers if the new version has any problems.

C) Eliminated - This approach introduces an extra Lambda function to filter requests, increasing complexity.

D) Eliminated - Creating a new API Gateway involves replicating configurations like authorizers, throttling, and endpoints. Maintaining two separate API Gateways adds long-term management overhead.

upvoted 1 times

🗨️ **rbryanben** 2 months, 3 weeks ago

Selected Answer: C

How not C

upvoted 1 times

🗨️ **nbxyzd** 3 months ago

Selected Answer: A

It's hands-down A.

upvoted 1 times

🗨️ **queekao** 6 months, 1 week ago

Selected Answer: A

the request is `Which solution will meet these requirements with the LEAST "operational overhead"?`, Therefore solution C is too crazy and the operation will be complex

upvoted 3 times

🗨️ **65703c1** 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 2 times

🗨️ 👤 **Alearn** 1 year, 1 month ago

Selected Answer: B

Option A requires creating a new stage on API Gateway, which might increase the operational overhead and complexity of managing multiple stages.

upvoted 1 times

🗨️ 👤 **NaghamAbdellatif** 1 year, 4 months ago

Why not B?

There is canary testing in Lambda Functions

upvoted 1 times

🗨️ 👤 **Cerakoted** 1 year, 3 months ago

Cuz of it -> new Lambda function for more testing by other developers with no impact to customers that use the API.

upvoted 3 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

Thank you for this... I too thought B --> definitely A then

upvoted 1 times

🗨️ 👤 **jayvarma** 1 year, 5 months ago

There is no need for us to create an all-new API Gateway in order to test the newer version of lambda. As a newer version of the lambda function is deployed with the necessary changes, a new stage of the API Gateway can be used to test the changes of the lambda function.

upvoted 3 times

🗨️ 👤 **jayvarma** 1 year, 5 months ago

So A is the right option

upvoted 2 times

🗨️ 👤 **jipark** 1 year, 6 months ago

Selected Answer: A

A : create new API stage (add stage) - correct

D: create new API Gateway (create new one) - incorrect

upvoted 4 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

yea D makes no sense. I think it was placed in there to throw people off.

upvoted 1 times

🗨️ 👤 **MrPie** 1 year, 6 months ago

Selected Answer: A

A is correct. Why the "correct answer" is always wrong? What's the point?

upvoted 4 times

🗨️ 👤 **JamalDaBoss** 1 year, 6 months ago

I agree, very stupid

upvoted 1 times

🗨️ 👤 **FunkyFresco** 1 year, 7 months ago

Selected Answer: A

A is ok according to my perspective.

upvoted 1 times

🗨️ 👤 **loctong** 1 year, 8 months ago

Selected Answer: A

A's true

upvoted 1 times

🗨️ 👤 **delak** 1 year, 8 months ago

Selected Answer: A

A is true



upvoted 1 times

🗨️ 👤 **rind2000** 1 year, 8 months ago

Selected Answer: A

In my perspective, A is the correct answer and a pretty typical pattern; I'm not sure why C was chosen, but testing in production is not a smart practice.

upvoted 1 times

  **chumji** 1 year, 8 months ago

The answer is A

upvoted 3 times

A company uses AWS Lambda functions and an Amazon S3 trigger to process images into an S3 bucket. A development team set up multiple environments in a single AWS account.


After a recent production deployment, the development team observed that the development S3 buckets invoked the production environment Lambda functions. These invocations caused unwanted execution of development S3 files by using production Lambda functions. The development team must prevent these invocations. The team must follow security best practices.

Which solution will meet these requirements?

- A. Update the Lambda execution role for the production Lambda function to add a policy that allows the execution role to read from only the production environment S3 bucket.
- B. Move the development and production environments into separate AWS accounts. Add a resource policy to each Lambda function to allow only S3 buckets that are within the same account to invoke the function.
- C. Add a resource policy to the production Lambda function to allow only the production environment S3 bucket to invoke the function.
- D. Move the development and production environments into separate AWS accounts. Update the Lambda execution role for each function to add a policy that allows the execution role to read from the S3 bucket that is within the same account.

Correct Answer: B

Community vote distribution



🗳️ 👤 **AgboolaKun** Highly Voted 1 year, 8 months ago

Selected Answer: C

B is a wrong answer because I do not understand the need to move the environments to separate AWS accounts. The resource policy in the production environment can be used to control which S3 bucket invokes the function.

In my understanding, the answer choice C fulfills the security best practices requirement in the question.

upvoted 27 times

🗳️ 👤 **MrPie** 1 year, 6 months ago

It's a best practice: Best Practices:

Separate workloads using accounts: Establish common guardrails and isolation between environments (such as production, development, and test) and workloads through a multi-account strategy. Account-level separation is strongly recommended, as it provides a strong isolation boundary for security, billing, and access. https://wa.aws.amazon.com/wat.question.SEC_1.en.html

upvoted 15 times

🗳️ 👤 **maurice2005** 10 months, 2 weeks ago

There is nowhere mentioned in the question that workload is the problem!

upvoted 1 times

🗳️ 👤 **wh1t4k3r** 5 months, 2 weeks ago

"The team must follow security best practices"

Security best practices state that prod should be separated from non prod environments.

upvoted 3 times

🗳️ 👤 **jipark** 1 year, 6 months ago

resource policy totally fulfill requirement

upvoted 3 times

🗳️ 👤 **csG13** Highly Voted 1 year, 7 months ago

Selected Answer: B

I choose B because it says that the team should follow the best security practices. AWS well-architected framework suggests separation. For reference see the link below: https://wa.aws.amazon.com/wat.question.SEC_1.en.html

upvoted 23 times

🗳️ 👤 **mooncake1** Most Recent 6 days, 22 hours ago



Selected Answer: C

Separating account can be more secure and easy but maintenance will be more hard.

It is developer's role and ability to distinguish and maintain between two environments, not just splitting them because it's easy.

What AWS wants as a DVA Certificate Member is not a person who splits the accounts because its easy , but is able to use IAM properly.

upvoted 1 times

  **rkotit** 1 week, 2 days ago

Selected Answer: C

When securing AWS Lambda functions, you must ensure that only authorized resources (like an S3 bucket) can trigger or invoke the function. In this scenario, the simplest and most effective solution is to use resource-based policies on the Lambda function to restrict access.

upvoted 1 times

  **sumanshu** 1 month ago



Selected Answer: C

A) Eliminated - This approach focuses on the Lambda function's ability to read from S3, not on which S3 buckets can invoke it.

B) Eliminated - Moving environments to separate AWS accounts might improve isolation but is not necessary to solve the stated problem

C) Correct - The resource policy in Option C restricts Lambda function invocations to only the intended production S3 bucket. This is a direct implementation of AWS's least privilege security model


upvoted 1 times

  **Dimix3** 1 month, 1 week ago

Selected Answer: C

Option C focuses on securing the production Lambda function by adding a resource policy to restrict invocations. This policy would specify that only the production S3 bucket can trigger the production Lambda function. This prevents accidental invocations from development S3 buckets.

upvoted 1 times

  **nbxyzd** 3 months ago

Selected Answer: B

Pls pay attention to keywords when you do aws quizzes.

The best security practice: option B

The lowest operational cost: option C.

So B is definitely the answer.

upvoted 2 times

  **AmitRanchi** 3 months, 3 weeks ago

Selected Answer: C

Because-

B. Move the development and production environments into separate AWS accounts:

While moving to separate accounts is a good practice for environment isolation, it's a more complex and potentially costly solution. Also, it isn't strictly necessary to achieve the goal of preventing unauthorized invocations, which can be accomplished via resource policies. It adds overhead without directly addressing the root cause.

upvoted 1 times

  **Saudis** 4 months, 2 weeks ago

Selected Answer: B



Best practices is the key Word

upvoted 1 times

  **Saurabh04** 5 months, 4 weeks ago

Option D is correct. This approach ensures isolation while maintaining manageability

upvoted 1 times

  **queekao** 6 months, 1 week ago

Selected Answer: C

Option C is the simplest way to achieve this requirement.

upvoted 1 times

  **nbxyzd** 3 months ago

Read carefully. The questions asks for 'the best' not the 'the simplest' security practice.

upvoted 1 times

🗨️ 👤 **Anandesh** 6 months, 1 week ago

Selected Answer: B

Establish common guardrails and isolation between environments (such as production, development, and test) and workloads through a multi-account strategy. Account-level separation is strongly recommended, as it provides a strong isolation boundary for security, billing, and access

upvoted 1 times

🗨️ 👤 **Aws_aspr** 6 months, 2 weeks ago

Asked 18 July 24 Without security best practices word. So C is correct answer for this.

upvoted 1 times

🗨️ 👤 **65703c1** 8 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗨️ 👤 **SerialIDr** 11 months, 1 week ago

Selected Answer: C

This approach involves configuring a resource-based policy (also known as a Lambda function policy) that explicitly defines which resources (in this case, S3 buckets) can invoke the Lambda function. By specifying only the production S3 bucket in the resource policy of the production Lambda function, you ensure that only events from the designated production S3 bucket can trigger the production Lambda function. This prevents development or other non-production buckets from inadvertently invoking production Lambda functions, thus maintaining environment integrity and security best practices.

upvoted 1 times

🗨️ 👤 **KarBiswa** 11 months, 1 week ago

Selected Answer: D

I feel it is D as there is no doubt we need to separately create two accounts for DEV & PROD. After that there must lambda execution roles where we can the specific policies. Resource based policies more of a Cross Account access.

<https://docs.aws.amazon.com/lambda/latest/dg/access-control-resource-based.html>

<https://repost.aws/knowledge-center/lambda-execution-role-s3-bucket>

As the question demands the best practices scenario so option D fulfils that.

upvoted 2 times

🗨️ 👤 **SD_CS** 11 months, 3 weeks ago

Selected Answer: B

I initially thought C, but after going through the below, I dont think there is any scope for doubt.

Establish common guardrails and isolation between environments (such as production, development, and test) and workloads through a multi-account strategy. Account-level separation is strongly recommended, as it provides a strong isolation boundary for security, billing, and access

https://docs.aws.amazon.com/en_us/wellarchitected/latest/framework/sec_securely_operate_multi_accounts.html

upvoted 4 times


A developer is creating an application. New users of the application must be able to create an account and register by using their own social media accounts.

Which AWS service or resource should the developer use to meet these requirements?

- A. IAM role
- B. Amazon Cognito identity pools
- C. Amazon Cognito user pools
- D. AWS Directory Service

Correct Answer: C

Community vote distribution



HuiHsin Highly Voted 1 year, 8 months ago

Selected Answer: C

<https://docs.aws.amazon.com/cognito/latest/developerguide/cognito-user-identity-pools.html>

upvoted 10 times

Bhatfield Highly Voted 1 year, 4 months ago

Amazon Cognito user pools provide user identity management and authentication for your application. They allow you to create and maintain a user directory, and you can enable social identity providers like Facebook, Google, or Amazon to allow users to register and log in using their social media accounts. This service is specifically designed for user management and authentication scenarios like the one described.

Option B, "Amazon Cognito identity pools," is more focused on providing temporary AWS credentials for users to access AWS services securely after they have been authenticated through a user pool.

upvoted 9 times

[Removed] 1 year, 1 month ago

The big difference being users authenticates to applications (web and mobile) vs identity authenticates to AWS resources.

upvoted 3 times

sumanshu Most Recent 1 month ago

Selected Answer: B

A) Eliminated - Eliminated because it does not support user account registration or social sign-in directly.

B) Correct - Manages user registration and authentication.

C) Eliminated - Provides temporary AWS credentials to authenticated users so they can access AWS resources like S3, DynamoDB, etc.

upvoted 1 times

sumanshu 1 month ago

Sorry - C

A) Eliminated - Eliminated because it does not support user account registration or social sign-in directly.

C) Correct - Manages user registration and authentication.

B) Eliminated - Provides temporary AWS credentials to authenticated users so they can access AWS resources like S3, DynamoDB, etc.

upvoted 1 times

nbxyzd 2 months, 4 weeks ago

Selected Answer: C

B is wrong. While identity pools also support federated identities, they are primarily used for granting AWS credentials to authenticated users. They are less focused on user management within the application itself.

upvoted 1 times

🗨️ 👤 **ogogundare** 3 months ago

The best answer is B Amazon Cognito Identity Pool

upvoted 1 times

🗨️ 👤 **ogogundare** 3 months ago

I think C is the right answer on further check

Amazon Cognito user pool can use social idps

upvoted 1 times

🗨️ 👤 **gdm83** 4 months, 3 weeks ago

Does emamtopic have exact replicas of actual exam questions

upvoted 1 times

🗨️ 👤 **65703c1** 8 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗨️ 👤 **65703c1** 8 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗨️ 👤 **SerialiDr** 11 months, 1 week ago

Selected Answer: B

Option B: Amazon Cognito identity pools

Amazon Cognito identity pools (also known as federated identities) enable you to create unique identities for your users and authenticate them with identity providers. With identity pools, your users can obtain temporary AWS credentials to access AWS services. This service supports authentication through social identity providers such as Amazon, Facebook, Google, and also supports unauthenticated identities.

upvoted 2 times

🗨️ 👤 **rrharris** 11 months, 2 weeks ago

c

Amazon Cognito user pools provide user identity management and authentication for your application.

upvoted 1 times

🗨️ 👤 **SerialiDr** 1 year ago

Selected Answer: B

B. Amazon Cognito identity pools: Amazon Cognito identity pools (also known as Federated Identities) enable you to create unique identities for your users and authenticate them with identity providers, including social media platforms like Facebook, Google, Amazon, and Apple. With identity pools, you can grant your users access to other AWS services. They are designed to handle scenarios where users can sign in through a third-party identity provider or use guest access.

upvoted 1 times

🗨️ 👤 **Dushank** 1 year, 4 months ago

Selected Answer: C

For creating an application where new users can create accounts and register using their social media accounts, Amazon Cognito is the most suitable service. Specifically, you'd want to use Amazon Cognito User Pools.

Amazon Cognito User Pools support sign-ins using social identity providers like Facebook, Google, and Amazon, as well as enterprise identity providers via SAML 2.0. With a user pool, you can create a fully managed user directory to enable user sign-up and sign-in, as well as handle password recovery, user verification, and other user management tasks.

upvoted 2 times

🗨️ 👤 **Dushank** 1 year, 4 months ago

The answer is (B).

Amazon Cognito identity pools is a managed service that provides user sign-in and identity management for your web and mobile applications. It supports social sign-in with a variety of providers, including Amazon, Facebook, Google, and Twitter.

upvoted 1 times

🗨️ 👤 **hanJR** 1 year, 9 months ago


Selected Answer: C

You can't register using Identity Pool. It lets you authenticate with provided identification pools.
upvoted 4 times

☒  **Cloud_Cloud** 1 year, 9 months ago

Selected Answer: C

<https://medium.com/wolox/integrating-social-media-to-your-app-with-aws-cognito-8943329aa89b>
upvoted 5 times


☒  **MrTee** 1 year, 9 months ago

Selected Answer: B

Key word is registration using their social media accounts
upvoted 4 times

☒  **awsdummie** 1 year, 8 months ago

B is incorrect. <https://www.youtube.com/watch?v=9pvvgKluCpl>
upvoted 1 times

☒  **rlnd2000** 1 year, 8 months ago

Using Cognito identity pools you can get the token and access AWS using social media accounts, BUT you can't create an account, in this case we need Cognito user pools.
upvoted 1 times

A social media application uses the AWS SDK for JavaScript on the frontend to get user credentials from AWS Security Token Service (AWS STS). The application stores its assets in an Amazon S3 bucket. The application serves its content by using an Amazon CloudFront distribution with the origin set to the S3 bucket.

The credentials for the role that the application assumes to make the SDK calls are stored in plaintext in a JSON file within the application code. The developer needs to implement a solution that will allow the application to get user credentials without having any credentials hardcoded in the application code.

Which solution will meet these requirements?

- A. Add a Lambda@Edge function to the distribution. Invoke the function on viewer request. Add permissions to the function's execution role to allow the function to access AWS STS. Move all SDK calls from the frontend into the function.
- B. Add a CloudFront function to the distribution. Invoke the function on viewer request. Add permissions to the function's execution role to allow the function to access AWS STS. Move all SDK calls from the frontend into the function.
- C. Add a Lambda@Edge function to the distribution. Invoke the function on viewer request. Move the credentials from the JSON file into the function. Move all SDK calls from the frontend into the function.
- D. Add a CloudFront function to the distribution. Invoke the function on viewer request. Move the credentials from the JSON file into the function. Move all SDK calls from the frontend into the function.

Correct Answer: A

Community vote distribution

A (88%)

13%

 **csG13** Highly Voted 1 year, 1 month ago

Selected Answer: A

The answer is A. Here is a reference directly from AWS docs:

"If you need some of the capabilities of Lambda@Edge that are not available with CloudFront Functions, such as network access or a longer execution time, you can still use Lambda@Edge before and after content is cached by CloudFront."

Since the requirement is to access the STS service, network access is required. Therefore, it can't be Cloudfront functions. Also, as a side note it's worth to mention that Cloudfront functions can only execute for up to 1ms. Apparently this isn't enough to fetch user creds (tokens) from STS.

The table in the following link summarises the differences between Cloudfront functions and Lambda@edge

<https://aws.amazon.com/blogs/aws/introducing-cloudfront-functions-run-your-code-at-the-edge-with-low-latency-at-any-scale/>
upvoted 14 times

 **MrTee** Highly Voted 1 year, 3 months ago

Selected Answer: B


The difference between A and B is the SDK for Javascript in use here; Lambda@Edge functions can be written in a variety of programming languages, including Node.js, Python, and Java, while CloudFront functions are written in JavaScript.

upvoted 5 times

 **Cloud_Cloud** 1 year, 3 months ago

Now one problem is lambda function can not perform AWS STS command

upvoted 1 times

 **eboehm** 1 year, 1 month ago

After rereading the last part of the question. It doesnt mention that it must remain written in Javascript, but does seem using AWS STS is a requirement so I think I would stick with A being the answer

upvoted 1 times

 **sumanshu** Most Recent 1 month ago

Selected Answer: A

B) Eliminated - CloudFront functions are lightweight JavaScript functions designed for simple HTTP request and response manipulations (e.g., header rewrites, URL rewrites). They cannot access AWS services like STS or assume roles

C/D) Eliminated - Moving the credentials from the JSON file to the Lambda@Edge function does not eliminate the core issue of hardcoding credentials

upvoted 1 times

🗨️ **65703c1** 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ **rrharris** 5 months, 2 weeks ago

Selected Answer: A

Why A is Correct:

Lambda@Edge for Secure Credential Management: Lambda@Edge allows you to run Lambda functions in response to CloudFront events. By using Lambda@Edge, the developer can securely manage credentials by keeping them out of the client-side code.

Invoking on Viewer Request: Invoking the Lambda@Edge function on viewer requests ensures that the credential generation happens in real-time, securely, and as needed, without exposing any sensitive information.

Execution Role with STS Access: Assigning the Lambda function an execution role with permissions to access AWS STS (Security Token Service) enables the function to securely request temporary, limited-privilege credentials on behalf of the client.

Moving SDK Calls to Lambda@Edge: Transferring all AWS SDK calls from the frontend to the Lambda@Edge function prevents exposing any credentials in the frontend code, enhancing security.

upvoted 3 times

🗨️ **SerialiDr** 6 months, 2 weeks ago

Selected Answer: A

A. Lambda@Edge allows you to run Lambda functions in response to CloudFront events. By using a Lambda@Edge function, you can securely handle the process of obtaining credentials from AWS STS without exposing them in the client-side application code. The function's execution role can be granted the necessary permissions to interact with AWS STS, and SDK calls can be made from within this server-side environment. This approach centralizes credential management and AWS interactions in a more secure, server-side context.

upvoted 4 times

🗨️ **LR2023** 8 months ago

I think i will also go with A as cloudfront functions can only read authorization headers from the viewer request if it sees the authorization header request. And Clouf front functions has no access to internet.

upvoted 2 times

🗨️ **Baba_Eni** 10 months, 4 weeks ago

Selected Answer: A

I will go for A, check the link below, Cloudfront functions are just within Cloudfront, hence, they DONT HAVE NETWORK ACCESS. Network access is required to make a call to AWS STS.

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/edge-functions.html>

upvoted 1 times

🗨️ **MG1407** 11 months, 3 weeks ago

The answer is B. I was in agreement with csG13 until a further research into the JavaScript SDK and STS. Found the following:

<https://docs.aws.amazon.com/AWSJavaScriptSDK/v3/latest/clients/client-cloudfront/classes/stsclient.html>.

Since the question states Js SDK and STS the answer is B.

upvoted 1 times

🗨️ **FunkyFresco** 1 year, 2 months ago

Selected Answer: A

Option A.


upvoted 1 times

🗨️ **zodraz** 1 year, 2 months ago

Selected Answer: A

<https://www.examttopics.com/discussions/amazon/view/89838-exam-aws-certified-developer-associate-topic-1-question-361/>

upvoted 2 times

 **vic614** 1 year, 3 months ago

Selected Answer: A

Cloud front function doesn't have network access, it has to be lambda @ edge

||

upvoted 2 times

An ecommerce website uses an AWS Lambda function and an Amazon RDS for MySQL database for an order fulfillment service. The service needs to return order confirmation immediately.


During a marketing campaign that caused an increase in the number of orders, the website's operations team noticed errors for "too many connections" from Amazon RDS. However, the RDS DB cluster metrics are healthy. CPU and memory capacity are still available.

What should a developer do to resolve the errors?

- A. Initialize the database connection outside the handler function. Increase the `max_user_connections` value on the parameter group of the DB cluster. Restart the DB cluster.
- B. Initialize the database connection outside the handler function. Use RDS Proxy instead of connecting directly to the DB cluster.
- C. Use Amazon Simple Queue Service (Amazon SQS) FIFO queues to queue the orders. Ingest the orders into the database. Set the Lambda function's concurrency to a value that equals the number of available database connections.
- D. Use Amazon Simple Queue Service (Amazon SQS) FIFO queues to queue the orders. Ingest the orders into the database. Set the Lambda function's concurrency to a value that is less than the number of available database connections.

Correct Answer: B

Community vote distribution



MrTee Highly Voted 1 year, 9 months ago

Selected Answer: B

Use an RDS Proxy instead of connecting directly to the DB cluster.

upvoted 13 times

hanJR Highly Voted 1 year, 9 months ago

B

<https://aws.amazon.com/blogs/compute/using-amazon-rds-proxy-with-aws-lambda/>

upvoted 7 times

sumanshu Most Recent 1 month ago

Selected Answer: B

The "too many connections" error occurs because Lambda functions create separate database connections for each invocation. RDS Proxy is designed to manage database connections efficiently by pooling and reusing connections, reducing the load on the database. Lambda functions can share connections from the proxy pool, avoiding the "too many connections" issue.

A) Eliminated - Increasing `max_user_connections` temporarily alleviates the issue but does not address the root cause

upvoted 1 times

Saudis 4 months, 2 weeks ago

Selected Answer: B

The key word is increase in the number of orders, so the proxy can help in this situation

upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

maurice2005 10 months, 2 weeks ago


Selected Answer: C

Option B only improve the execution time of lambda and decrease the delay from request to database. It might even worsen the situation because database can get more concurrent connection. RDS Proxy also doesn't limit the number of connections, even if so it will generate errors for lambda.

The only way is to throttle the requests using SQS until a connection gets released.

Why everyone thinks ChatGPT has the ultimate answer????!!!!

upvoted 2 times

  **nbxyzd** 2 months, 4 weeks ago

Read carefully: "the service needs to return order confirmation *immediately*."

This means that asynchronous solutions like SQS does NOT meet the requirement.



upvoted 2 times

  **SerialiDr** 1 year ago

Selected Answer: B

AWS RDS Proxy is designed to manage and pool database connections, which makes it ideal for environments with highly variable and potentially high-volume database access patterns, such as those driven by Lambda functions. It helps to reduce the number of direct connections to the database and can efficiently manage the connections from the pool.

upvoted 3 times

  **hmdev** 1 year, 5 months ago

Selected Answer: B

We can use an RDS proxy to handle a lot of connections. We are choosing this option because the load on the RDS is normal. If the RDS was unable to handle loads, we would've checked other options like queues or transactions.



upvoted 2 times

  **eberhe900** 1 year, 6 months ago

Selected Answer: B

<https://repost.aws/questions/QLXSqEPGbQx6qiyBa1D1Udg/lambda-to-db-connectivity-best-practices>

upvoted 1 times

  **loctong** 1 year, 8 months ago

Selected Answer: B

Using an RDS Proxy can manage connections to the RDS instance, reducing the overhead of establishing new connections and thereby preventing the "too many connections" error.

upvoted 2 times

A company stores its data in data tables in a series of Amazon S3 buckets. The company received an alert that customer credit card information might have been exposed in a data table on one of the company's public applications. A developer needs to identify all potential exposures within the application environment.

Which solution will meet these requirements?

- A. Use Amazon Athena to run a job on the S3 buckets that contain the affected data. Filter the findings by using the SensitiveData:S3Object/Personal finding type.
- B. Use Amazon Macie to run a job on the S3 buckets that contain the affected data. Filter the findings by using the SensitiveData:S3Object/Financial finding type.
- C. Use Amazon Macie to run a job on the S3 buckets that contain the affected data. Filter the findings by using the SensitiveData:S3Object/Personal finding type.
- D. Use Amazon Athena to run a job on the S3 buckets that contain the affected data. Filter the findings by using the SensitiveData:S3Object/Financial finding type.

Correct Answer: B

Community vote distribution

B (100%)

MrTee **Highly Voted** 1 year, 9 months ago

Selected Answer: B

Use Amazon Macie to run a job on the S3 buckets that contain the affected data. Filter the findings by using the SensitiveData:S3Object/Financial finding type.

Option A and D suggest using Amazon Athena, which is an interactive query service that can be used to analyze data stored in S3 using standard SQL queries. While Athena can help identify data in S3 buckets, it does not provide the same level of automated scanning and pattern matching that Amazon Macie does.

Option C is incorrect because the SensitiveData:S3Object/Personal finding type is designed to identify personally identifiable information (PII), such as names and addresses, but not credit card information.

upvoted 23 times

sumanshu **Most Recent** 1 month ago

Selected Answer: B

Macie is a fully managed data security and privacy service that uses machine learning and pattern matching to discover and protect sensitive data stored in Amazon S3.

C) Eliminated - SensitiveData:S3Object/Personal finding type is irrelevant because credit card information is not classified as "personal" in this context but as "financial."

A/D) Eliminated - Athena does not support sensitive data detection

upvoted 1 times

Saudis 4 months, 2 weeks ago

Selected Answer: B

The Ans is B because the credit card is a sensitive Data and also it is a financial Data
upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

SD_CS 12 months ago

Selected Answer: B

SensitiveData:S3Object/Financial only works with Macie?? so how can it be D?

upvoted 1 times

🗨️ 👤 **SerialiDr** 1 year ago

Selected Answer: B

B. Use Amazon Macie to run a job on the S3 buckets that contain the affected data. Filter the findings by using the SensitiveData:S3Object/Financial finding type: Amazon Macie is a security service that uses machine learning and pattern matching to discover and protect sensitive data in AWS. Macie is designed to identify various types of sensitive data, including financial data, which would cover credit card information. This option is suitable for the requirement as it leverages Macie's capability to specifically identify and report on exposures of sensitive financial data.

upvoted 3 times

🗨️ 👤 **Baba_Eni** 1 year, 7 months ago

Selected Answer: B

<https://docs.aws.amazon.com/macie/latest/user/findings-types.html>

upvoted 4 times

🗨️ 👤 **HuiHsin** 1 year, 7 months ago

Selected Answer: B

https://docs.aws.amazon.com/zh_tw/macie/latest/user/findings-types.html

upvoted 1 times

🗨️ 👤 **Prem28** 1 year, 8 months ago

Selected Answer: B

The best solution to identify all potential exposures within the application environment after receiving an alert that customer credit card information might have been exposed in a data table on one of the company's public applications is to use Amazon Macie. Amazon Macie is a fully managed data security and privacy service that uses machine learning and pattern matching to discover and protect sensitive data in AWS.

upvoted 1 times

A software company is launching a multimedia application. The application will allow guest users to access sample content before the users decide if they want to create an account to gain full access. The company wants to implement an authentication process that can identify users who have already created an account. The company also needs to keep track of the number of guest users who eventually create an account.

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

- A. Create an Amazon Cognito user pool. Configure the user pool to allow unauthenticated users. Exchange user tokens for temporary credentials that allow authenticated users to assume a role.
- B. Create an Amazon Cognito identity pool. Configure the identity pool to allow unauthenticated users. Exchange unique identity for temporary credentials that allow all users to assume a role.
- C. Create an Amazon CloudFront distribution. Configure the distribution to allow unauthenticated users. Exchange user tokens for temporary credentials that allow all users to assume a role.
- D. Create a role for authenticated users that allows access to all content. Create a role for unauthenticated users that allows access to only the sample content.
- E. Allow all users to access the sample content by default. Create a role for authenticated users that allows access to the other content.

Correct Answer: BD

Community vote distribution

BD (100%)

MrTee Highly Voted 1 year, 3 months ago

Selected Answer: BD

option B because by configuring the identity pool to allow unauthenticated users, you can enable guest users to access the sample content. When users create an account, they can be authenticated, and then given access to the full content by assuming a role that allows them access. Option D is correct because creating roles for authenticated and unauthenticated users with different levels of access is an appropriate way to meet the requirement of identifying users who have created an account and keeping track of the number of guest users who eventually create an account.

upvoted 25 times

sumanshu Most Recent 1 month ago

Selected Answer: BD

A) Eliminated - User pools are primarily for managing authenticated users. They do not natively support unauthenticated (guest) users.

B) Identity pools support both unauthenticated (guest) and authenticated users

C) Eliminated - CloudFront is a content delivery service, not an authentication or identity management service.

D) This allows proper segregation of access, where guest users can only access sample content while authenticated users gain full access.

E) Eliminated - Proper identity management requires differentiating between roles for guest and authenticated users.

upvoted 1 times

65703c1 2 months, 1 week ago

Selected Answer: BD

BD is the correct answer.

upvoted 1 times

a_win 7 months, 1 week ago

Selected Answer: BD

E won't be a choice because "The company also needs to keep track of the number of guest users who eventually create an account."

upvoted 4 times

KarBiswa 7 months, 2 weeks ago

Selected Answer: BD

Covers Unauthenticated and authenticated users scenario

upvoted 1 times

 **jjpark** 12 months ago

Selected Answer: BD

"who already created account" means User Pool not required. - NOT A

upvoted 4 times

A company is updating an application to move the backend of the application from Amazon EC2 instances to a serverless model. The application uses an Amazon RDS for MySQL DB instance and runs in a single VPC on AWS. The application and the DB instance are deployed in a private subnet in the VPC.

The company needs to connect AWS Lambda functions to the DB instance.

Which solution will meet these requirements?

- A. Create Lambda functions inside the VPC with the AWSLambdaBasicExecutionRole policy attached to the Lambda execution role. Modify the RDS security group to allow inbound access from the Lambda security group.
- B. Create Lambda functions inside the VPC with the AWSLambdaVPCLambdaAccessExecutionRole policy attached to the Lambda execution role. Modify the RDS security group to allow inbound access from the Lambda security group.
- C. Create Lambda functions with the AWSLambdaBasicExecutionRole policy attached to the Lambda execution role. Create an interface VPC endpoint for the Lambda functions. Configure the interface endpoint policy to allow the lambda:InvokeFunction action for each Lambda function's Amazon Resource Name (ARN).
- D. Create Lambda functions with the AWSLambdaVPCLambdaAccessExecutionRole policy attached to the Lambda execution role. Create an interface VPC endpoint for the Lambda functions. Configure the interface endpoint policy to allow the lambda:InvokeFunction action for each Lambda function's Amazon Resource Name (ARN).

Correct Answer: B

Community vote distribution

B (81%)

D (19%)

 **MrTee** Highly Voted 1 year, 3 months ago

Selected Answer: B

The AWSLambdaVPCLambdaAccessExecutionRole policy allows the Lambda function to create elastic network interfaces (ENIs) in the VPC and use the security groups attached to those ENIs for controlling inbound and outbound traffic.

upvoted 14 times

 **Khaja2k** Most Recent 1 month ago

Selected Answer: B

B: This option ensures that the Lambda functions have the necessary permissions to access resources within the VPC and that the RDS security group is configured to allow inbound access from the Lambda functions.

D: Reason for eliminating this is creating an interface VPC endpoint and configuring the endpoint policy is unnecessary for this scenario as it complicates the setup without addressing the VPC access requirements directly.

upvoted 1 times

 **sumanshu** 1 month ago

Selected Answer: B

Lambda functions need to run inside the VPC to access resources like the RDS instance, which is located within the private subnet.

The AWSLambdaVPCLambdaAccessExecutionRole policy is required for Lambda functions to connect to resources inside a VPC. This policy allows Lambda functions to use Elastic Network Interfaces (ENIs) to connect to the VPC.


The security group associated with the RDS instance must allow inbound connections from the Lambda function's security group.

upvoted 2 times

 **sumanshu** 1 month ago

A) Eliminated - AWSLambdaBasicExecutionRole provides basic permissions to write logs to Amazon CloudWatch but does not grant the necessary permissions for Lambda to connect to resources in a VPC

upvoted 1 times

 **65703c1** 2 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗨️ **SerialiDr** 6 months, 2 weeks ago

Selected Answer: B

This is the correct solution. The AWSLambdaVPCLambdaAccessExecutionRole policy includes permissions that allow the Lambda function to access resources within a VPC, such as an RDS instance. Additionally, modifying the RDS security group to allow inbound access from the Lambda security group is necessary to enable network connectivity between the Lambda functions and the RDS instance.

upvoted 3 times

🗨️ **KarBiswa** 7 months, 2 weeks ago

Selected Answer: B

<https://docs.aws.amazon.com/lambda/latest/dg/lambda-intro-execution-role.html>

upvoted 1 times

🗨️ **Nagasoacle** 9 months, 2 weeks ago

Selected Answer: D

Answer : D

upvoted 2 times

🗨️ **love777** 11 months, 1 week ago

Selected Answer: D

While Lambda functions cannot run directly in private subnets, they can be configured to access resources within a VPC by creating a VPC endpoint for Lambda.

AWS Lambda supports VPC Endpoints for Lambda, which allow Lambda functions to securely access resources within a VPC without needing to traverse the public internet.

You should attach the AWSLambdaVPCLambdaAccessExecutionRole policy to your Lambda execution role to enable it to create network interfaces in your VPC for accessing resources.

By configuring an interface VPC endpoint for Lambda, you can enable the Lambda function to communicate with resources within the private subnet and the RDS instance.

upvoted 4 times

🗨️ **Baba_Eni** 1 year, 1 month ago

Selected Answer: B

<https://docs.aws.amazon.com/aws-managed-policy/latest/reference/AWSLambdaVPCLambdaAccessExecutionRole.html>

<https://docs.aws.amazon.com/lambda/latest/dg/lambda-intro-execution-role.html>

upvoted 3 times

🗨️ **Prem28** 1 year, 1 month ago

ans- opt d

Option A does not allow Lambda functions to access resources in the VPC.

Option B does not create an interface VPC endpoint, which means that Lambda functions will be exposed to the public internet.

Option C does not configure the interface endpoint policy to allow the lambda:InvokeFunction action, which means that Lambda functions will not be able to invoke each other.

upvoted 3 times

🗨️ **jipark** 12 months ago

I definitely agree.

Lambda cannot be installed inside VPC, instead, AWSLambdaVPCLambdaAccessExecutionRole allow to connect via ENI.

upvoted 1 times

A company has a web application that runs on Amazon EC2 instances with a custom Amazon Machine Image (AMI). The company uses AWS CloudFormation to provision the application. The application runs in the us-east-1 Region, and the company needs to deploy the application to the us-west-1 Region.

An attempt to create the AWS CloudFormation stack in us-west-1 fails. An error message states that the AMI ID does not exist. A developer must resolve this error with a solution that uses the least amount of operational overhead.

Which solution meets these requirements?

- A. Change the AWS CloudFormation templates for us-east-1 and us-west-1 to use an AWS AMI. Relaunch the stack for both Regions.
- B. Copy the custom AMI from us-east-1 to us-west-1. Update the AWS CloudFormation template for us-west-1 to refer to AMI ID for the copied AMI. Relaunch the stack.
- C. Build the custom AMI in us-west-1. Create a new AWS CloudFormation template to launch the stack in us-west-1 with the new AMI ID.
- D. Manually deploy the application outside AWS CloudFormation in us-west-1.

Correct Answer: B

Community vote distribution

B (100%)

MrTee **Highly Voted** 1 year, 9 months ago

Selected Answer: B

This will allow the company to deploy the application to the us-west-1 Region using the same custom AMI that is used in the us-east-1 Region.
upvoted 12 times

sumanshu **Most Recent** 1 month ago

Selected Answer: B

A) Eliminated - The goal is to deploy the custom AMI, not AWS-provided AMI

C) Eliminated - Building the custom AMI from scratch in us-west-1 is unnecessary and adds extra operational overhead

D) Eliminated - it bypasses CloudFormation's infrastructure-as-code benefits and increases operational overhead.
upvoted 1 times

Saudis 4 months, 2 weeks ago

The Ans is B the keyword is operational overhead so from the choice the last overhead is copy
upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: B

B is the correct answer.
upvoted 1 times

gomurali 1 year, 7 months ago

<https://www.examttopics.com/discussions/amazon/view/78848-exam-aws-certified-developer-associate-topic-1-question-118/>
upvoted 2 times

A developer is updating several AWS Lambda functions and notices that all the Lambda functions share the same custom libraries. The developer wants to centralize all the libraries, update the libraries in a convenient way, and keep the libraries versioned.

Which solution will meet these requirements with the LEAST development effort?

- A. Create an AWS CodeArtifact repository that contains all the custom libraries.
- B. Create a custom container image for the Lambda functions to save all the custom libraries.
- C. Create a Lambda layer that contains all the custom libraries.
- D. Create an Amazon Elastic File System (Amazon EFS) file system to store all the custom libraries.

Correct Answer: C

Community vote distribution

C (100%)

MrTee **Highly Voted** 1 year, 9 months ago

Selected Answer: C

the most efficient solution is to use a Lambda layer to store the common libraries, update them in one place, and reference them from each Lambda function that requires them.

upvoted 18 times

sumanshu **Most Recent** 1 month ago

Selected Answer: C

Lambda layers allow you to manage your shared libraries separately from the function code

A) Eliminated - it would require additional steps for Lambda functions to access these libraries. The Lambda functions would need to fetch the libraries at runtime, adding more complexity to your setup.

B/D) Additional overhead

upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

HuiHsin 1 year, 7 months ago

Selected Answer: C

The Lambda layer of option C provides a simpler solution without the need to introduce an additional CodeArtifact service.

upvoted 3 times

loctong 1 year, 8 months ago

Selected Answer: C

Lambda layers are a distribution mechanism for libraries, custom runtimes, and other function dependencies in AWS Lambda. By creating a Lambda layer, you can package and centrally manage the shared custom libraries for the Lambda functions.

upvoted 2 times

loctong 1 year, 8 months ago

Selected Answer: C

It should be Create a Lambda layer.

upvoted 1 times

Ryan1002 1 year, 8 months ago

Why not CodeArtifact?

"CodeArtifact allows you to store artifacts using popular package managers and build tools like Maven, Gradle, npm, Yarn, Twine, pip, and NuGet.

CodeArtifact can automatically fetch software packages on demand from public package repositories so you can access the latest versions of application dependencies."

upvoted 2 times

🗨️ 👤 **nbxyzd** 2 months, 4 weeks ago

"LEAST development effort"

While CodeArtifact can be integrated with Lambda, it typically involves additional steps such as setting up credentials and configuring the Lambda functions to access the repository. This can require more development effort compared to Lambda Layer which is specifically designed for Lambda.

upvoted 1 times

🗨️ 👤 **jipark** 1 year, 6 months ago

"LEAST development effort"

upvoted 3 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

We are updating a Lambda function. Lambda layers are specifically used for situations mentioned in this question

upvoted 1 times

A developer wants to use AWS Elastic Beanstalk to test a new version of an application in a test environment.

Which deployment method offers the FASTEST deployment?

- A. Immutable
- B. Rolling
- C. Rolling with additional batch
- D. All at once

Correct Answer: D

Community vote distribution

D (100%)

  **yeacuz** Highly Voted  1 year, 2 months ago

Selected Answer: D

The answer is D.

"All at once – The quickest deployment method." <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.deploy-existing-version.html>

upvoted 8 times

  **loctong** Highly Voted  1 year, 2 months ago

Selected Answer: D

The "All at once" deployment method deploys the new version of the application to all instances simultaneously. It updates all instances of the environment in a short period of time, resulting in the fastest overall deployment.

upvoted 6 times

  **sumanshu** Most Recent  1 month ago

Selected Answer: D



This deployment method deploys the new version of the application to all instances in the environment simultaneously

A) Eliminated - In an immutable deployment, Elastic Beanstalk creates a new environment with the new version of the application and swaps it with the old one once it's ready. This method ensures that there is no impact on the current environment during the deployment, but it takes longer due to the creation of a new environment and the time needed for verification.

B) Eliminated - In a rolling deployment, Elastic Beanstalk updates a batch of instances one at a time while keeping the others running. This method minimizes downtime but takes longer than "all at once" since it updates instances in batches.

C) Eliminated -

upvoted 1 times

  **65703c1** 2 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

A company is providing read access to objects in an Amazon S3 bucket for different customers. The company uses IAM permissions to restrict access to the S3 bucket. The customers can access only their own files.

Due to a regulation requirement, the company needs to enforce encryption in transit for interactions with Amazon S3.

Which solution will meet these requirements?

- A. Add a bucket policy to the S3 bucket to deny S3 actions when the `aws:SecureTransport` condition is equal to false.
- B. Add a bucket policy to the S3 bucket to deny S3 actions when the `s3:x-amz-acl` condition is equal to `public-read`.
- C. Add an IAM policy to the IAM users to enforce the usage of the AWS SDK.
- D. Add an IAM policy to the IAM users that allows S3 actions when the `s3:x-amz-acl` condition is equal to `bucket-owner-read`.

Correct Answer: A

Community vote distribution

A (100%)

MrTee **Highly Voted** 1 year, 3 months ago

Selected Answer: A

This solution enforces encryption in transit for interactions with Amazon S3 by denying access to the S3 bucket if the request is not made over an HTTPS connection. This condition can be enforced by using the `"aws:SecureTransport"` condition key in a bucket policy.

upvoted 19 times

jipark 12 months ago

'in transit' = SSL Secure Transport

upvoted 2 times

sumanshu **Most Recent** 1 month ago

Selected Answer: A

By adding a bucket policy that denies S3 actions when the `aws:SecureTransport` condition is false, you ensure that only requests made over HTTPS are allowed. The condition `aws:SecureTransport` checks if the request was made using a secure transport (HTTPS).

B) Eliminated - `s3:x-amz-acl` condition relates to access control lists (ACLs) for the objects

C) Enforcing the usage of the AWS SDK does not guarantee encryption in transit.

D) Eliminated - The `bucket-owner-read` ACL allows the bucket owner to read the objects, but it does not enforce secure communication

upvoted 1 times

65703c1 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

loctong 1 year, 2 months ago

Selected Answer: A

To enforce encryption in transit for interactions with Amazon S3, you can add a bucket policy to the S3 bucket that denies S3 actions when the `aws:SecureTransport` condition is equal to false. This condition checks whether the requests to S3 are made over a secure (HTTPS) connection.

upvoted 3 times

rInd2000 1 year, 2 months ago

Selected Answer: A

<https://repost.aws/knowledge-center/s3-bucket-policy-for-config-rule>

upvoted 3 times

A company has an image storage web application that runs on AWS. The company hosts the application on Amazon EC2 instances in an Auto Scaling group. The Auto Scaling group acts as the target group for an Application Load Balancer (ALB) and uses an Amazon S3 bucket to store the images for sale.

The company wants to develop a feature to test system requests. The feature will direct requests to a separate target group that hosts a new beta version of the application.

Which solution will meet this requirement with the LEAST effort?

- A. Create a new Auto Scaling group and target group for the beta version of the application. Update the ALB routing rule with a condition that looks for a cookie named version that has a value of beta. Update the test system code to use this cookie to test the beta version of the application.
- B. Create a new ALB, Auto Scaling group, and target group for the beta version of the application. Configure an alternate Amazon Route 53 record for the new ALB endpoint. Use the alternate Route 53 endpoint in the test system requests to test the beta version of the application.
- C. Create a new ALB, Auto Scaling group, and target group for the beta version of the application. Use Amazon CloudFront with Lambda@Edge to determine which specific request will go to the new ALB. Use the CloudFront endpoint to send the test system requests to test the beta version of the application.
- D. Create a new Auto Scaling group and target group for the beta version of the application. Update the ALB routing rule with a condition that looks for a cookie named version that has a value of beta. Use Amazon CloudFront with Lambda@Edge to update the test system requests to add the required cookie when the requests go to the ALB.

Correct Answer: A

Community vote distribution



MrTee Highly Voted 1 year, 9 months ago

Selected Answer: A

This solution will allow the company to direct requests to a separate target group that hosts the new beta version of the application without having to create a new ALB or use additional services such as Amazon Route 53 or Amazon CloudFront.

Option D adds additional complexity and effort compared to option A, which simply involves updating the ALB routing rule with a condition that looks for a cookie named version that has a value of beta and updating the test system code to use this cookie to test the beta version of the application.

upvoted 26 times

backfringe Highly Voted 1 year, 6 months ago

Selected Answer: B

Option B provides the simplest and least effort solution to test the beta version of the application. By creating a new ALB, Auto Scaling group, and target group for the beta version, the company can deploy the new version of the application separately from the production version.

Configuring an alternate Amazon Route 53 record for the new ALB endpoint allows the test system requests to be directed to the beta version.

upvoted 8 times

wh1t4k3r 5 months, 2 weeks ago

I will say that it is way simpler to create an ALB rule to a new tg and alter the beta app then create a new tg, a new alb and configure route 53 and then deploy the app

upvoted 1 times

sumanshu Most Recent 1 month ago

Selected Answer: A

B) Eliminated - You need a second load balancer, which increases both costs and maintenance overhead.

Testers must switch between different URLs to test the beta version, which is less seamless.

C/D) Eliminated - Lambda@Edge increases complexity unnecessarily

upvoted 1 times

🗨️ 👤 **sumanshu** 1 month ago

A)

New Beta Group:

You create a new "group of servers" (Auto Scaling group + target group) for the beta version of your app. These servers will run the updated beta code.

Routing with Cookies:

The existing load balancer is updated to check incoming requests for a specific cookie in the browser.

If the cookie says something like version=beta, the load balancer will route the request to the beta servers.

If there's no cookie, the load balancer routes the request to the normal servers (current version).

Test System's Role:

The test system (used by testers) adds the special cookie (version=beta) to its requests. This ensures only the testers access the beta version. Regular users don't have this cookie, so they continue to see the normal version.

upvoted 1 times

🗨️ 👤 **Anandesh** 6 months, 3 weeks ago

Selected Answer: A

<https://aws.amazon.com/blogs/aws/new-advanced-request-routing-for-aws-application-load-balancers/>

upvoted 1 times

🗨️ 👤 **65703c1** 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ 👤 **lvRa** 9 months, 3 weeks ago

ChatGPT goes for D.

upvoted 1 times

🗨️ 👤 **41eb566** 10 months, 2 weeks ago

Selected Answer: B

The solution that will meet the requirement with the least effort is:

B. Create a new ALB, Auto Scaling group, and target group for the beta version of the application. Configure an alternate Amazon Route 53 record for the new ALB endpoint. Use the alternate Route 53 endpoint in the test system requests to test the beta version of the application.

upvoted 1 times

🗨️ 👤 **SerialiDr** 11 months, 1 week ago

Selected Answer: A

This approach allows for the least amount of effort in setting up a beta environment where test system requests can be directed to a new version of the application for testing purposes. It leverages ALB's ability to conditionally route traffic based on request attributes, such as cookies, allowing for flexible and efficient testing of new application versions alongside existing production workloads.

upvoted 1 times

🗨️ 👤 **SerialiDr** 1 year ago

Selected Answer: A

A. Create a new Auto Scaling group and target group for the beta version of the application. Update the ALB routing rule with a condition that looks for a cookie named version that has a value of beta. Update the test system code to use this cookie to test the beta version of the application: This is a straightforward and effective solution. By creating a new Auto Scaling group and target group for the beta version and updating the ALB to route based on a specific cookie, the company can easily direct test traffic to the beta version without needing additional infrastructure or complex configurations. The test system would simply include the specified cookie in its requests to access the beta version.

upvoted 1 times

🗨️ 👤 **JohnPI** 1 year ago

Selected Answer: D

A is modifying the code for testing, not a good practice. D is the least effort compared to B and C

upvoted 2 times

🗨️ 👤 **gqs3119** 1 year, 1 month ago

Selected Answer: D

Modifying ALB (A/D) is less effort than modifying route 53 and adding ALB (B/C), 1 action vs 2.

So it's A or D, let's think about effort in both cases.

In case of A you will need to:

1. Add a new temporary code to set cookies
2. Test app with new temporary code, to make sure it won't break the production
3. Deploy it to the production

After tests are finished:

4. Remove temporary code
5. Deploy to production

In case of D you will need:



1. Create lambda
2. Do a simple testing to make sure it won't affect production

After tests are finished:

3. Remove lambda

I'd say D is the least effort.

upvoted 3 times

  **a_win** 1 year, 1 month ago

Selected Answer: A

requirement with the LEAST effort

upvoted 1 times

  **LR2023** 1 year, 1 month ago

Selected Answer: D

just using voting...explanation in a different thread

upvoted 3 times

  **LR2023** 1 year, 1 month ago

I am going with D.....

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-at-the-edge.html>

A Lambda function can inspect cookies and rewrite URLs so that users see different versions of a site for A/B testing.

Option B & C requires to create new ALB - which is not least effort. And option A requires to update code.

upvoted 3 times

  **Nagadoracle** 1 year, 3 months ago

Selected Answer: B

Considering Least effort


upvoted 3 times

  **LemonGremlin** 1 year, 3 months ago

Selected Answer: A

Agree that this is A

upvoted 1 times

  **Rameez1** 1 year, 3 months ago

Selected Answer: A

Option A serves the requirement with least efforts.

upvoted 1 times


A team is developing an application that is deployed on Amazon EC2 instances. During testing, the team receives an error. The EC2 instances are unable to access an Amazon S3 bucket.

Which steps should the team take to troubleshoot this issue? (Choose two.)

- A. Check whether the policy that is assigned to the IAM role that is attached to the EC2 instances grants access to Amazon S3.
- B. Check the S3 bucket policy to validate the access permissions for the S3 bucket.
- C. Check whether the policy that is assigned to the IAM user that is attached to the EC2 instances grants access to Amazon S3.
- D. Check the S3 Lifecycle policy to validate the permissions that are assigned to the S3 bucket.
- E. Check the security groups that are assigned to the EC2 instances. Make sure that a rule is not blocking the access to Amazon S3.

Correct Answer: AB

Community vote distribution



MrTee Highly Voted 1 year, 3 months ago

Selected Answer: AB

Option A is correct because IAM roles are used to grant permissions to AWS services, such as EC2 instances, to access other AWS services, such as S3 buckets. The policy assigned to the IAM role attached to the EC2 instances should be checked to ensure that it grants access to the S3 bucket.

Option B is also correct because the S3 bucket policy controls access to the S3 bucket. The S3 bucket policy should be checked to ensure that the access permissions are correctly configured.

upvoted 17 times

sumanshu Most Recent 1 month ago

Selected Answer: AB

EC2 instances typically assume an IAM role to interact with AWS services like S3. If the attached IAM role does not have the correct permissions, the EC2 instance will not be able to access the S3 bucket.

Even if the IAM role grants S3 permissions, the bucket policy might explicitly deny access or restrict access to certain principals (users or roles).

C) Eliminated - EC2 instances do not use IAM users to access AWS services.

D) Eliminated - An S3 Lifecycle policy is used for managing object lifecycles (e.g., moving objects to Glacier or deleting old versions). It does not control access permissions.

E) Eliminated - Security groups control inbound and outbound traffic at the network level

upvoted 1 times

65703c1 2 months, 1 week ago

Selected Answer: AB

AB is the correct answer.

upvoted 1 times

65703c1 2 months, 1 week ago

AB is the correct answer.

upvoted 1 times

konieczny69 6 months ago

Selected Answer: AB

Incorrectly stated question. Its not mentioned how does the application us IAM, that is wether its STS or user credentials. AC is as well perfectly correct answer.

upvoted 1 times

SerialIDr 6 months, 2 weeks ago

Selected Answer: AB

The two steps most relevant to troubleshooting the issue are:

- A. Check whether the policy that is assigned to the IAM role that is attached to the EC2 instances grants access to Amazon S3.
- B. Check the S3 bucket policy to validate the access permissions for the S3 bucket.

upvoted 2 times

Nagasaoracle 9 months, 2 weeks ago

Selected Answer: AB

<https://repost.aws/knowledge-center/ec2-instance-access-s3-bucket>

upvoted 4 times

love777 11 months, 1 week ago

Selected Answer: AE

Explanation:

A. IAM Role Policy: EC2 instances are typically associated with IAM roles. These roles have policies attached to them that define the permissions the instances have. If the instances are unable to access an S3 bucket, it's essential to verify that the IAM role assigned to the EC2 instances has the necessary permissions to interact with S3.

E. Security Groups: Security groups act as virtual firewalls for EC2 instances. They control inbound and outbound traffic. If the EC2 instances are unable to access S3, it's possible that the associated security group is blocking outbound traffic to the S3 service. Make sure the security group rules allow outbound traffic to the S3 service.

upvoted 3 times

love777 11 months, 1 week ago

The correct steps to troubleshoot the issue are:

- A. Check whether the policy that is assigned to the IAM role that is attached to the EC2 instances grants access to Amazon S3.
- E. Check the security groups that are assigned to the EC2 instances. Make sure that a rule is not blocking the access to Amazon S3.

Explanation:

E. Security Groups: Security groups act as virtual firewalls for EC2 instances. They control inbound and outbound traffic. If the EC2 instances are unable to access S3, it's possible that the associated security group is blocking outbound traffic to the S3 service. Make sure the security group rules allow outbound traffic to the S3 service.

upvoted 2 times

awsazedevesh 1 year ago

Why not E ?

upvoted 2 times

remynick 11 months, 3 weeks ago

access to S3 is controlled by IAM, not security groups.

upvoted 3 times

maurice2005 4 months, 2 weeks ago

Security group is like a firewall, can block any inbound/outbound traffic.

upvoted 1 times

indirasubbaraj 1 year, 1 month ago

AB

<https://repost.aws/knowledge-center/ec2-instance-access-s3-bucket>

upvoted 1 times

Prem28 1 year, 1 month ago

AE



B. Check the S3 bucket policy to validate the access permissions for the S3 bucket. The S3 bucket policy controls who has access to the bucket, but it does not control how they can access it. The IAM role or user that is attached to the EC2 instances must have the appropriate permissions to access the bucket, regardless of what the S3 bucket policy says.

C. Check whether the policy that is assigned to the IAM user that is attached to the EC2 instances grants access to Amazon S3. This is unlikely to

be the cause of the issue, as the IAM role is what is typically used to control access to AWS resources.

D. Check the S3 Lifecycle policy to validate the permissions that are assigned to the S3 bucket. The S3 Lifecycle policy controls how objects are stored and moved in Amazon S3. It does not control who has access to the bucket.

upvoted 2 times

  **vic614** 1 year, 3 months ago

Selected Answer: AB

A: Make sure EC2 instance profile has permission to access s3

B: Make sure S3 resource policy allows the access from instance

upvoted 4 times

A developer is working on an ecommerce website. The developer wants to review server logs without logging in to each of the application servers individually. The website runs on multiple Amazon EC2 instances, is written in Python, and needs to be highly available.

How can the developer update the application to meet these requirements with MINIMUM changes?

- A. Rewrite the application to be cloud native and to run on AWS Lambda, where the logs can be reviewed in Amazon CloudWatch.
- B. Set up centralized logging by using Amazon OpenSearch Service, Logstash, and OpenSearch Dashboards.
- C. Scale down the application to one larger EC2 instance where only one instance is recording logs.
- D. Install the unified Amazon CloudWatch agent on the EC2 instances. Configure the agent to push the application logs to CloudWatch.

Correct Answer: D

Community vote distribution

D (100%)

MrTee **Highly Voted** 1 year, 3 months ago

Selected Answer: D

Option D is the best option because it requires minimum changes and leverages the existing infrastructure.

upvoted 12 times

sumanshu **Most Recent** 1 month ago

Selected Answer: D

By installing the unified Amazon CloudWatch agent on all EC2 instances, the application logs will automatically be pushed to Amazon CloudWatch Logs, a centralized logging service.

upvoted 1 times

65703c1 2 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

SerialiDr 6 months, 2 weeks ago

Selected Answer: D

D. Install the unified Amazon CloudWatch agent on the EC2 instances. Configure the agent to push the application logs to CloudWatch: This is the most appropriate solution. The unified CloudWatch agent can be easily installed and configured on each EC2 instance to push logs to Amazon CloudWatch. This allows for centralized log storage and access without a significant change to the application architecture or its high availability setup. It provides a straightforward way to aggregate logs from multiple instances in one place.

upvoted 3 times

loctong 1 year, 2 months ago

Selected Answer: D

By installing the Amazon CloudWatch agent on the EC2 instances, the developer can easily collect and send logs from each instance to Amazon CloudWatch. The CloudWatch agent provides a unified way to collect logs, system-level metrics, and custom metrics from the EC2 instances.

upvoted 3 times

A company is creating an application that processes .csv files from Amazon S3. A developer has created an S3 bucket. The developer has also created an AWS Lambda function to process the .csv files from the S3 bucket.

Which combination of steps will invoke the Lambda function when a .csv file is uploaded to Amazon S3? (Choose two.)

- A. Create an Amazon EventBridge rule. Configure the rule with a pattern to match the S3 object created event.
- B. Schedule an Amazon EventBridge rule to run a new Lambda function to scan the S3 bucket.
- C. Add a trigger to the existing Lambda function. Set the trigger type to EventBridge. Select the Amazon EventBridge rule.
- D. Create a new Lambda function to scan the S3 bucket for recently added S3 objects.
- E. Add S3 Lifecycle rules to invoke the existing Lambda function.

Correct Answer: AC

Community vote distribution

AC (92%)

4%

MrTee **Highly Voted** 1 year, 9 months ago

Selected Answer: AC

Option A is correct because an Amazon EventBridge rule can be created to detect when an object is created in an S3 bucket. The rule should be configured with a pattern to match the S3 object created event.

Option C is correct because the existing Lambda function can be updated with an EventBridge trigger. The trigger type should be set to EventBridge, and the Amazon EventBridge rule created in step A should be selected.

upvoted 17 times

sumanshu **Most Recent** 1 month ago

Selected Answer: AC

EventBridge allows you to trigger actions (like Lambda functions) based on specific events (like S3 object uploads).

upvoted 1 times

wh1t4k3r 5 months, 2 weeks ago

Selected Answer: AC

For the options given as answers, A and C is the viable option. But its a stupid solution... you can use s3 events to trigger it directly.

upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: AC

AC is the correct answer.

upvoted 1 times

tqiu654 1 year, 1 month ago

Selected Answer: AE

ChatGPT:AE

upvoted 1 times

Hari4455 1 year, 1 month ago

ChatGPT: AC

A. Create an Amazon EventBridge rule. Configure the rule with a pattern to match the S3 object created event.

This sets up an EventBridge rule to respond to S3 object creation events.

C. Add a trigger to the existing Lambda function. Set the trigger type to EventBridge. Select the Amazon EventBridge rule.

This associates the Lambda function with the EventBridge rule, ensuring that the Lambda function is triggered when the specified event occurs.

upvoted 1 times

tqiu654 1 year, 1 month ago

Lambda functions are not currently supported as triggers directly from EventBridge rules. Lambda can be used as the target of an EventBridge rule, but is not added to a Lambda function as a trigger.

upvoted 1 times

🗨️ 👤 **Nagasaoracle** 1 year, 3 months ago

Selected Answer: AC

AC is combination of steps required

upvoted 2 times

🗨️ 👤 **Jing2023** 1 year, 3 months ago

Why not just use the S3 event as the trigger directly.

upvoted 3 times

🗨️ 👤 **ValeriiRadchenko** 1 year, 2 months ago

I agree that in general this is a stupid question. But maybe company need's to use EB in application 🤔

upvoted 1 times

🗨️ 👤 **Naj_64** 1 year, 6 months ago

Selected Answer: AC

A C for sure

upvoted 2 times

🗨️ 👤 **loctong** 1 year, 8 months ago

Selected Answer: AB

A, B are correctly

upvoted 1 times

A developer needs to build an AWS CloudFormation template that self-populates the AWS Region variable that deploys the CloudFormation template.

What is the MOST operationally efficient way to determine the Region in which the template is being deployed?

- A. Use the AWS::Region pseudo parameter.
- B. Require the Region as a CloudFormation parameter.
- C. Find the Region from the AWS::StackId pseudo parameter by using the Fn::Split intrinsic function.
- D. Dynamically import the Region by referencing the relevant parameter in AWS Systems Manager Parameter Store.

Correct Answer: A

Community vote distribution

A (100%)

MrTee **Highly Voted** 1 year, 9 months ago

Selected Answer: A

A. Use the AWS::Region pseudo parameter.
upvoted 10 times

sumanshu **Most Recent** 1 month ago

Selected Answer: A

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/mappings-section-structure.html>

AWS::Region is a pseudo parameter that automatically provides the Region where the stack is being deployed (e.g., us-east-1 if deploying in the US East Region).

B) Eliminated - Adding a parameter for the Region means users would need to manually specify the Region every time they deploy the template

C) Eliminated - While the stack ID contains the Region as part of its value (e.g., arn:aws:cloudformation:us-east-1:123456789012:stack/MyStack/...), extracting it with Fn::Split is a complicated workaround.

D) Eliminated
upvoted 1 times

albert_kuo 4 months, 1 week ago

Selected Answer: A

Resources:

MyS3Bucket:

Type: "AWS::S3::Bucket"

Properties:

BucketName: !Sub "my-bucket-\${AWS::Region}"

upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

Baba_Eni 1 year, 7 months ago

Selected Answer: A

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


upvoted 2 times

Baba_Eni 1 year, 7 months ago

Selected Answer: A

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

upvoted 1 times

 **loctong** 1 year, 8 months ago

Selected Answer: A

The `AWS::Region` pseudo parameter is a built-in CloudFormation parameter that automatically resolves to the AWS Region where the CloudFormation stack is being created. By using this pseudo parameter, you can dynamically access the current Region without requiring any additional configuration or input.

upvoted 4 times

A company has hundreds of AWS Lambda functions that the company's QA team needs to test by using the Lambda function URLs. A developer needs to configure the authentication of the Lambda functions to allow access so that the QA IAM group can invoke the Lambda functions by using the public URLs.

Which solution will meet these requirements?

- A. Create a CLI script that loops on the Lambda functions to add a Lambda function URL with the `AWS_IAM` auth type. Run another script to create an IAM identity-based policy that allows the `lambda:InvokeFunctionUrl` action to all the Lambda function Amazon Resource Names (ARNs). Attach the policy to the QA IAM group.
- B. Create a CLI script that loops on the Lambda functions to add a Lambda function URL with the `NONE` auth type. Run another script to create an IAM resource-based policy that allows the `lambda:InvokeFunctionUrl` action to all the Lambda function Amazon Resource Names (ARNs). Attach the policy to the QA IAM group.
- C. Create a CLI script that loops on the Lambda functions to add a Lambda function URL with the `AWS_IAM` auth type. Run another script to loop on the Lambda functions to create an IAM identity-based policy that allows the `lambda:InvokeFunctionUrl` action from the QA IAM group's Amazon Resource Name (ARN).
- D. Create a CLI script that loops on the Lambda functions to add a Lambda function URL with the `NONE` auth type. Run another script to loop on the Lambda functions to create an IAM resource-based policy that allows the `lambda:InvokeFunctionUrl` action from the QA IAM group's Amazon Resource Name (ARN).

Correct Answer: A

Community vote distribution

A (74%)

C (26%)

MrTee **Highly Voted** 1 year, 9 months ago

Selected Answer: A

Option A meets these requirements?

upvoted 16 times

ppardav 1 year, 7 months ago

<https://docs.aws.amazon.com/lambda/latest/dg/urls-auth.html>

upvoted 4 times

jipark 1 year, 6 months ago

create 'AWS_IAM auth type' -> Attach the policy to the QA IAM group

upvoted 5 times

sumanshu **Most Recent** 1 month ago

Selected Answer: A

A) Correct - Setting the `AWS_IAM` authentication type ensures that only IAM users or roles with the right permissions can invoke the Lambda function URLs.

C) Eliminated - While this approach is secure, creating individual resource-based policies for hundreds of Lambda functions is unnecessarily complex and hard to manage. An identity-based policy (used in Option A) is simpler because it applies to the entire QA IAM group at once.

B/D - Eliminated - Setting `NONE` as the auth type makes the Lambda function URLs publicly accessible without authentication.

upvoted 1 times

Anandesh 6 months, 3 weeks ago

Selected Answer: A

Apologies again, please refer to the youtube link I shared earlier..correct ans is A

upvoted 1 times

Anandesh 7 months ago

I think the answer is B here, reason being the function should be invoked using public urls

upvoted 1 times

🗨️ **65703c1** 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ **SerialIDr** 11 months, 1 week ago

Selected Answer: A

This approach leverages AWS IAM authentication (AWS_IAM auth type) for Lambda function URLs, ensuring that only authenticated and authorized IAM entities can invoke the Lambda functions. By creating an IAM policy that specifies the `lambda:InvokeFunctionUrl` action and attaching it to the QA IAM group, you provide the necessary permissions for the QA team to invoke the Lambda functions securely. This method aligns with AWS best practices for security and access control, allowing for scalable and manageable access management across multiple Lambda functions.

upvoted 1 times

🗨️ **CrescentShared** 11 months, 1 week ago

Selected Answer: C

I don't know why so much A, but isn't A giving the access to all the lambda?

upvoted 1 times

🗨️ **SD_CS** 11 months, 3 weeks ago

Selected Answer: A

I have to go for A even though it appears both should suffice. I took this from AWS Documentation

If you choose the `AWS_IAM` auth type, users who need to invoke your Lambda function URL must have the `lambda:InvokeFunctionUrl` permission. Depending on who makes the invocation request, you may have to grant this permission using a resource-based policy.

If the principal making the request is in the same AWS account as the function URL, then the principal must either have `lambda:InvokeFunctionUrl` permissions in their identity-based policy, OR have permissions granted to them in the function's resource-based policy.

AWS clearly states both should be good. The reason for selecting A is the wording is clear, loop on to lambda function to provide the permission was bit of confusing to me.

upvoted 1 times

🗨️ **konieczny69** 12 months ago

Selected Answer: C

I don't get all A answers. This is typical resource based policy that allows invoking a function by concrete principal - in this case its the QA role.

For all those who vote for A - go ahead and create simple API Gateway with a lambda integration type. Then look at the resource based policy - `lambda:InvokeFunction` allowed by `apigateway.amazonaws.com` with `ArnLike` condition.

ChatGTP also says C.

upvoted 2 times

🗨️ **love777** 1 year, 5 months ago

Selected Answer: C

Explanation:

In this scenario, the QA team needs to test AWS Lambda functions using Lambda function URLs while ensuring proper authentication and access control. Here's why option C is the appropriate solution:

Authentication Type: Using the `AWS_IAM` auth type for the Lambda function URLs ensures that the Lambda functions can be invoked only by users and roles that have the necessary IAM permissions.

Identity-Based Policy: By creating an IAM identity-based policy, you grant permissions directly to the QA IAM group to invoke the Lambda functions using the Lambda function URLs. This provides fine-grained control over which IAM entities can access the functions.

Option A uses the `AWS_IAM` auth type and creates a policy for the QA IAM group, which is a good direction. However, the creation of a policy that allows `lambda:InvokeFunctionUrl` for all Lambda function ARNs might grant excessive permissions.

upvoted 4 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

pay attention to the wording of the answers:

A - Run another script to create an IAM identity-based policy that allows the lambda:InvokeFunctionUrl action to all the Lambda function Amazon Resource Names (ARNs).

*This option is very clear. You are creating an IAM identity based policy allowing access to invoke the function and then attaching this policy to the QA IAM group.

C - Run another script to loop on the Lambda functions to create an IAM identity-based policy that allows the lambda:InvokeFunctionUrl action from the QA IAM group's Amazon Resource Name (ARN).

*What does "Run another script to loop on the Lambda functions" What does this even mean?? are we doing some sort of while loop here?

Wording for this option is very confusing and makes no sense to me. I go with A

upvoted 4 times

🗨️ 👤 **Manel87** 1 year, 1 month ago

good thought!

upvoted 1 times

🗨️ 👤 **dezoito** 1 year, 3 months ago

Why A grant excessive permissions? The policy will contain only the Lambda's ARNs wich the QA group should have access to.

upvoted 2 times

A developer maintains a critical business application that uses Amazon DynamoDB as the primary data store. The DynamoDB table contains millions of documents and receives 30-60 requests each minute. The developer needs to perform processing in near-real time on the documents when they are added or updated in the DynamoDB table.

How can the developer implement this feature with the LEAST amount of change to the existing application code?

- A. Set up a cron job on an Amazon EC2 instance. Run a script every hour to query the table for changes and process the documents.
- B. Enable a DynamoDB stream on the table. Invoke an AWS Lambda function to process the documents.
- C. Update the application to send a PutEvents request to Amazon EventBridge. Create an EventBridge rule to invoke an AWS Lambda function to process the documents.
- D. Update the application to synchronously process the documents directly after the DynamoDB write.

Correct Answer: B

Community vote distribution

B (100%)

MrTee **Highly Voted** 1 year, 3 months ago

Selected Answer: B

Option B is the best solution because it proposes enabling a DynamoDB stream on the table, which allows the developer to capture document-level changes in near-real time without modifying the application code. Then, the stream can be configured to invoke an AWS Lambda function to process the documents in near-real time. This solution requires minimal changes to the existing application code, and the Lambda function can be developed and deployed separately, enabling the developer to easily maintain and update it as needed.

upvoted 10 times

loctong **Highly Voted** 1 year, 2 months ago

Selected Answer: B

To implement near-real-time processing on documents added or updated in a DynamoDB table with the least amount of change to the existing application code, the developer should:

B. Enable a DynamoDB stream on the table and invoke an AWS Lambda function to process the documents.

Enabling a DynamoDB stream on the table allows capturing and processing of the changes made to the table in near-real-time. The stream provides an ordered sequence of item-level modifications (inserts, updates, and deletes) that can be consumed by other AWS services, such as AWS Lambda.

upvoted 5 times

sumanshu **Most Recent** 1 month ago

Selected Answer: B

A DynamoDB stream is a feature that automatically tracks changes (inserts, updates, deletes) to items in a table. Whenever a document is added or updated, the change is recorded in the stream.

A) Eliminated - Querying the DynamoDB table every hour introduces a delay, so it's not near-real-time.

C) Eliminated - To use EventBridge, you would need to modify the application code to explicitly send events (PutEvents) every time a change is made in DynamoDB.

D) Eliminated - require significant changes to the existing application code.

upvoted 1 times

65703c1 2 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

SerialiDr 6 months, 2 weeks ago

Selected Answer: B

GPT

To implement near-real-time processing of documents when they are added or updated in an Amazon DynamoDB table with the least amount of change to the existing application code, let's evaluate the options:

A. Set up a cron job on an Amazon EC2 instance. Run a script every hour to query the table for changes and process the documents: This approach introduces additional complexity and is not near-real time. Running a script periodically to check for updates is inefficient and does not meet the requirement for immediate processing upon document addition or update.

B. Enable a DynamoDB stream on the table. Invoke an AWS Lambda function to process the documents: This is the most efficient and least intrusive option. DynamoDB Streams capture changes to items in the DynamoDB table as they occur in near-real time and can trigger an AWS Lambda function automatically. This setup requires minimal changes to the existing application code, as the processing logic is moved to the Lambda function, which is triggered by the stream events.

upvoted 1 times

A developer is writing an application for a company. The application will be deployed on Amazon EC2 and will use an Amazon RDS for Microsoft SQL Server database. The company's security team requires that database credentials are rotated at least weekly.

How should the developer configure the database credentials for this application?

- A. Create a database user. Store the user name and password in an AWS Systems Manager Parameter Store secure string parameter. Enable rotation of the AWS Key Management Service (AWS KMS) key that is used to encrypt the parameter.
- B. Enable IAM authentication for the database. Create a database user for use with IAM authentication. Enable password rotation.
- C. Create a database user. Store the user name and password in an AWS Secrets Manager secret that has daily rotation enabled.
- D. Use the EC2 user data to create a database user. Provide the user name and password in environment variables to the application.

Correct Answer: C

Community vote distribution

C (100%)

MrTee **Highly Voted** 1 year, 9 months ago

Selected Answer: C

option C: Create a database user. Store the user name and password in an AWS Secrets Manager secret that has daily rotation enabled. This will allow the developer to securely store the database credentials and automatically rotate them at least weekly to meet the company's security requirements.

upvoted 14 times

sumanshu **Most Recent** 1 month ago

Selected Answer: C

It can automatically rotate the credentials for supported databases

A) Eliminated - it does not natively support automatic credential rotation for databases.

upvoted 1 times

Saurabh04 5 months, 3 weeks ago

Correct Answer should be A: This approach centralizes credential management and provides secure storage. Rotation can be scheduled weekly as required by the security team1.

upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

SerialiDr 1 year ago

Selected Answer: C

C. Create a database user. Store the user name and password in an AWS Secrets Manager secret that has daily rotation enabled: This is the correct solution. AWS Secrets Manager is specifically designed to handle secrets like database credentials, including their rotation. You can configure Secrets Manager to automatically rotate the credentials as frequently as needed (e.g., daily or weekly), which aligns with the security team's requirements.

upvoted 2 times

jjpark 1 year, 6 months ago

Selected Answer: C

rotation key & cross account key is feature of Secret Manager

<https://tutorialsdojo.com/aws-secrets-manager-vs-systems-manager-parameter-store/>

upvoted 2 times

Baba_Eni 1 year, 7 months ago



Selected Answer: C

<https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html>

https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotate-secrets_turn-on-for-other.html

https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotate-secrets_schedule.html

upvoted 3 times

  **loctong** 1 year, 8 months ago

Selected Answer: C

the keyword is "rotation"

upvoted 4 times


A real-time messaging application uses Amazon API Gateway WebSocket APIs with backend HTTP service. A developer needs to build a feature in the application to identify a client that keeps connecting to and disconnecting from the WebSocket connection. The developer also needs the ability to remove the client.

Which combination of changes should the developer make to the application to meet these requirements? (Choose two.)

- A. Switch to HTTP APIs in the backend service.
- B. Switch to REST APIs in the backend service.
- C. Use the callback URL to disconnect the client from the backend service.
- D. Add code to track the client status in Amazon ElastiCache in the backend service.
- E. Implement \$connect and \$disconnect routes in the backend service.

Correct Answer: DE

Community vote distribution



MrTee Highly Voted 1 year, 9 months ago

Selected Answer: DE

Option D because by storing the client status in the cache, the backend service can quickly access the client status data without the need to query the database or perform other time-consuming operations.

Option E. Implement \$connect and \$disconnect routes in the backend service: \$connect and \$disconnect are the reserved routes in WebSocket APIs, which are automatically called by API Gateway whenever a client connects or disconnects from the WebSocket. By implementing these routes in the backend service, the developer can track and manage the client status, including identifying and removing the client when needed.

upvoted 22 times

akmv2 10 months, 1 week ago

How is D viable when the question doesn't mention ElastiCache? You're making an assumption that adding a service/configuration is what's being asked

upvoted 1 times

catcatpunch Highly Voted 1 year, 8 months ago

Selected Answer: CE

C => https://docs.aws.amazon.com/ko_kr/apigateway/latest/developerguide/apigateway-how-to-call-websocket-api-connections.html

E => https://docs.aws.amazon.com/ko_kr/apigateway/latest/developerguide/apigateway-websocket-api-route-keys-connect-disconnect.html

upvoted 10 times

mooncake1 Most Recent 6 days, 12 hours ago

Selected Answer: CE

DE does not make the developer disconnect

upvoted 1 times

mooncake1 6 days, 12 hours ago

Can't believe DE is 50%.. this q is not even hard

upvoted 1 times

Arad 3 weeks, 4 days ago

Selected Answer: CE

CE is correct.

upvoted 1 times

sumanshu 1 month ago

Selected Answer: DE

A) Eliminated - Switching to HTTP APIs would eliminate the WebSocket capability entirely,

B) Eliminated - REST APIs are also unsuitable for WebSocket connections, as they are designed for request-response interactions rather than persistent, real-time communication.

Both C&E and D&E - seems correct

C) - API Gateway provides a callback URL that your backend service can use to invoke the POST /connections/{connectionId} API with the disconnect command.

D)

E) Correct - The \$connect route triggers when a client establishes a WebSocket connection. The \$disconnect route triggers when a client disconnects from the WebSocket

upvoted 1 times

🗨️ 👤 **sumanshu** 1 month ago

C) Callback URL allows you to explicitly disconnect clients.

upvoted 1 times

🗨️ 👤 **sumanshu** 1 month ago

The two requirements in the question are:

Identify a client that keeps connecting and disconnecting:

This requires tracking client connections and disconnections over time.

To achieve this, the backend needs a mechanism to log or store connection events for each client, which can then be analyzed to identify repeated connection patterns.

Remove the client:

Once the client is identified as problematic (e.g., frequent connects/disconnects), the backend needs to remove the client. This is achieved using the WebSocket callback URL to explicitly disconnect the client.

upvoted 1 times

🗨️ 👤 **sumanshu** 1 month ago

D is critical for tracking client activity.

E is required to capture connection and disconnection events.

C provides the functionality to disconnect a client once identified.

Thus, the correct combination is C, D, and E.

upvoted 1 times

🗨️ 👤 **sumanshu** 1 month ago

C. Use the callback URL to disconnect the client from the backend service

This is essential for removing the client. Without it, the backend has no way to forcibly disconnect a problematic client.

☐ Must-have.

D. Add code to track the client status in Amazon ElastiCache in the backend service

Tracking client connection/disconnection behavior is critical to identifying problematic clients. ElastiCache is a highly suitable tool for storing connection data efficiently.

☐ Must-have.

E. Implement \$connect and \$disconnect routes in the backend service

While these routes are necessary for detecting client connections and disconnections, this could be seen as implicit in the implementation of D (since tracking client status assumes these routes are already used).

△ Optional, depending on how the question is interpreted.

upvoted 1 times

🗨️ 👤 **Saurabh04** 5 months, 3 weeks ago

Option DE: Implement \$connect and \$disconnect Routes:

Add \$connect and \$disconnect routes to your WebSocket API.

These routes handle client connections and disconnections.

When a client connects (\$connect), a Lambda function can add the connection ID to a data store (e.g., DynamoDB).

When a client disconnects (\$disconnect), another Lambda function can remove the connection ID from the data store.

Track Client Status:

Use Amazon ElastiCache (e.g., Redis) to track client status.
Store relevant information (e.g., client IDs, connection timestamps) in ElastiCache.
This allows you to identify clients that connect and disconnect
upvoted 1 times

🗨️ **9d8dd9c** 3 months, 1 week ago
What about "the ability to remove the client"? for that it needs to use the callback function
upvoted 1 times

🗨️ **tomchandler077** 6 months, 2 weeks ago
Option D --> CORRECT Because, tracking the client's connection status using ElastiCache could help in identifying clients with erratic connection patterns and managing stateful information in a distributed environment, which is useful for WebSocket applications.

Also Option E CORRECT. These routes handle connection and disconnection events.
upvoted 1 times

🗨️ **tsangckl** 7 months, 2 weeks ago
This appear at 17 Jun exam
upvoted 2 times

🗨️ **frangesk** 5 months, 1 week ago
Do you remember the answers?
upvoted 1 times

🗨️ **65703c1** 8 months, 1 week ago
Selected Answer: DE
DE is the correct answer.
upvoted 1 times

🗨️ **41eb566** 10 months, 2 weeks ago
Selected Answer: DE
E. Implement \$connect and \$disconnect routes in the backend service.

By implementing \$connect and \$disconnect routes, the backend service can capture when clients connect and disconnect from the WebSocket connection. This allows the application to track client status effectively.
D. Add code to track the client status in Amazon ElastiCache in the backend service.
upvoted 1 times

🗨️ **maurice2005** 10 months, 3 weeks ago
Selected Answer: CE
Your backend service can use the following WebSocket connection HTTP requests to send a callback message to a connected client, get connection information, or disconnect the client

<https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-how-to-call-websocket-api-connections.html>
upvoted 1 times

🗨️ **SerialiDr** 11 months, 1 week ago
Selected Answer: DE
When a client connects to your WebSocket API, the \$connect route is invoked, and when they disconnect, the \$disconnect route is invoked. You can use these routes to track the state of each client. By maintaining a record of each client's connections and disconnections, possibly in a database or an in-memory data store like Amazon ElastiCache, you can identify clients that frequently connect and disconnect.

Hence, the combination of changes that should be made to the application to meet these requirements includes:

Implement \$connect and \$disconnect routes in the backend service (Option E).
Add code to track the client status in Amazon ElastiCache in the backend service (Option D).
upvoted 2 times

🗨️ **KarBiswa** 11 months, 1 week ago
Selected Answer: CE
C option - [Supports https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-how-to-call-websocket-api-connections.html](https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-how-to-call-websocket-api-connections.html)
E option supports - <https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-websocket-api-route-keys-connect-disconnect.html>

upvoted 1 times

🗨️ 👤 **Ashwinvm22** 1 year ago

Selected Answer: CD

C: https://docs.aws.amazon.com/ko_kr/apigateway/latest/developerguide/apigateway-how-to-call-websocket-api-connections.html

D: You need a way to track which user is continuously reconnecting. That is why option D is so important because without it you will just be disconnecting every user that tries to connect cause then how will you know which user is the "problem" user. Note that you don't need the \$disconnect endpoint to disconnect a client if you use option C. So CD is the only combination to solve the problem.

upvoted 2 times

🗨️ 👤 **Abdullah22** 1 year ago

going with DE

upvoted 1 times

🗨️ 👤 **SerialiDr** 1 year ago

Selected Answer: CD

C. Use the callback URL to disconnect the client from the backend service: The callback URL can be used to send messages to connected clients or to disconnect them from the WebSocket connection. This approach allows the backend service to programmatically disconnect a client, which is useful for managing clients that frequently connect and disconnect.

D. Add code to track the client status in Amazon ElastiCache in the backend service: Implementing client status tracking in the backend service, possibly using a fast, in-memory data store like Amazon ElastiCache, allows the application to monitor and record the behavior of each client. This can be used to identify clients with frequent connect/disconnect patterns.

upvoted 1 times

🗨️ 👤 **a_win** 1 year, 1 month ago

Selected Answer: DE

D. Add code to track the client status in Amazon ElastiCache in the backend service.

E. Implement \$connect and \$disconnect routes in the backend service.

upvoted 1 times

A developer has written code for an application and wants to share it with other developers on the team to receive feedback. The shared application code needs to be stored long-term with multiple versions and batch change tracking.

Which AWS service should the developer use?

- A. AWS CodeBuild
- B. Amazon S3
- C. AWS CodeCommit
- D. AWS Cloud9

Correct Answer: C

Community vote distribution

C (100%)

MrTee **Highly Voted** 1 year, 3 months ago

Selected Answer: C

option C, AWS CodeCommit.

upvoted 7 times

sumanshu **Most Recent** 1 month ago

Selected Answer: C

AWS CodeCommit is specifically built for managing source code.

upvoted 1 times

65703c1 2 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

LR2023 8 months ago

Selected Answer: C

Code commit is a code source repository

upvoted 1 times

loctong 1 year, 2 months ago

Selected Answer: C

must be C

upvoted 3 times

delak 1 year, 2 months ago

it's C

upvoted 2 times

A company's developer is building a static website to be deployed in Amazon S3 for a production environment. The website integrates with an Amazon Aurora PostgreSQL database by using an AWS Lambda function. The website that is deployed to production will use a Lambda alias that points to a specific version of the Lambda function.

The company must rotate the database credentials every 2 weeks. Lambda functions that the company deployed previously must be able to use the most recent credentials.

Which solution will meet these requirements?

- A. Store the database credentials in AWS Secrets Manager. Turn on rotation. Write code in the Lambda function to retrieve the credentials from Secrets Manager.
- B. Include the database credentials as part of the Lambda function code. Update the credentials periodically and deploy the new Lambda function.
- C. Use Lambda environment variables. Update the environment variables when new credentials are available.
- D. Store the database credentials in AWS Systems Manager Parameter Store. Turn on rotation. Write code in the Lambda function to retrieve the credentials from Systems Manager Parameter Store.

Correct Answer: A

Community vote distribution

A (100%)

MrTee **Highly Voted** 1 year, 3 months ago

Selected Answer: A

Option A is the correct solution; Option D is also a valid solution, but it is not the best option since Secrets Manager provides built-in rotation, which ensures that the latest credentials are automatically updated. Additionally, AWS Systems Manager Parameter Store does not provide the ability to rotate secrets automatically.

upvoted 12 times

loctong **Highly Voted** 1 year, 2 months ago

Selected Answer: A

the key word is "rotation"

upvoted 5 times

sumanshu **Most Recent** 1 month ago

Selected Answer: A

AWS Secrets Manager provides a built-in feature for rotating credentials automatically

D) Eliminated - it does not provide built-in rotation functionality

upvoted 1 times

65703c1 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

SerialiDr 6 months, 2 weeks ago

Selected Answer: A

A. Store the database credentials in AWS Secrets Manager. Turn on rotation. Write code in the Lambda function to retrieve the credentials from Secrets Manager: This is the most suitable solution. AWS Secrets Manager is designed to manage, retrieve, and rotate secrets such as database credentials. By storing the credentials in Secrets Manager and enabling rotation, the credentials will be automatically rotated every 2 weeks. The Lambda function can retrieve the latest credentials programmatically from Secrets Manager, ensuring it always has access to the current credentials.

upvoted 2 times

LR2023 8 months ago

Selected Answer: A

Secrets manager for rotation
upvoted 1 times

A developer is developing an application that uses signed requests (Signature Version 4) to call other AWS services. The developer has created a canonical request, has created the string to sign, and has calculated signing information.

Which methods could the developer use to complete a signed request? (Choose two.)

- A. Add the signature to an HTTP header that is named Authorization.
- B. Add the signature to a session cookie.
- C. Add the signature to an HTTP header that is named Authentication.
- D. Add the signature to a query string parameter that is named X-Amz-Signature.
- E. Add the signature to an HTTP header that is named WWW-Authenticate.

Correct Answer: AD

Community vote distribution

AD (100%)

  **vicvega** Highly Voted 1 year ago

Header:

```
Authorization: AWS4-HMAC-SHA256
Credential=AKIAIOSFODNN7EXAMPLE/20220830/us-east-1/ec2/aws4_request,
SignedHeaders=host;x-amz-date,
Signature=calculated-signature
```

Query String:

```
https://ec2.amazonaws.com/?
Action=DescribeInstances&
Version=2016-11-15&
X-Amz-Signature=calculated-signature
```

<https://docs.aws.amazon.com/IAM/latest/UserGuide/create-signed-request.html>
upvoted 10 times

  **MrTee** Highly Voted 1 year, 3 months ago

Selected Answer: AD

the correct options are A and D.

upvoted 9 times

  **sumanshu** Most Recent 1 month ago



Selected Answer: AD

The developer has already performed the necessary steps to generate the signature but needs to know where to place the signature in the request to ensure the request is valid and properly authenticated.

A) Correct - SigV4 officially supports placing the signature in the Authorization header,

C) Eliminated - AWS specifically requires the Authorization header for passing the signature in HTTP headers

D) Correct - SigV4 allows for query string authentication as an alternative to using the Authorization header. This is especially useful for presigned URLs, where the signature and other necessary information are passed in the query string (e.g., X-Amz-Signature)
upvoted 1 times

  **65703c1** 2 months, 1 week ago

Selected Answer: AD

AD is the correct answer.

upvoted 1 times

🗨️ 👤 **SerialiDr** 6 months, 2 weeks ago

Selected Answer: AD

A. Add the signature to an HTTP header that is named Authorization: This is a correct method. In Signature Version 4, the completed signature is typically added to the request's Authorization header. This header includes the signing information along with other necessary components such as the Credential Scope and the Signed Headers.

D. Add the signature to a query string parameter that is named X-Amz-Signature: This is a correct method. In addition to including the signature in the Authorization header, Signature Version 4 also allows for presigned URLs where the signature is part of the query string parameters. The signature is included in the X-Amz-Signature query string parameter.

upvoted 3 times

🗨️ 👤 **loctong** 1 year, 2 months ago

Selected Answer: AD

Option B,C And E are not correct;

upvoted 1 times

🗨️ 👤 **awsdummie** 1 year, 2 months ago

Selected Answer: AD

<https://docs.aws.amazon.com/IAM/latest/UserGuide/create-signed-request.html>

upvoted 2 times

A company must deploy all its Amazon RDS DB instances by using AWS CloudFormation templates as part of AWS CodePipeline continuous integration and continuous delivery (CI/CD) automation. The primary password for the DB instance must be automatically generated as part of the deployment process.

Which solution will meet these requirements with the LEAST development effort?

- A. Create an AWS Lambda-backed CloudFormation custom resource. Write Lambda code that generates a secure string. Return the value of the secure string as a data field of the custom resource response object. Use the CloudFormation Fn::GetAtt intrinsic function to get the value of the secure string. Use the value to create the DB instance.
- B. Use the AWS CodeBuild action of CodePipeline to generate a secure string by using the following AWS CLI command: `aws secretsmanager get-random-password`. Pass the generated secure string as a CloudFormation parameter with the NoEcho attribute set to true. Use the parameter reference to create the DB instance.
- C. Create an AWS Lambda-backed CloudFormation custom resource. Write Lambda code that generates a secure string. Return the value of the secure string as a data field of the custom resource response object. Use the CloudFormation Fn::GetAtt intrinsic function to get a value of the secure string. Create secrets in AWS Secrets Manager. Use the secretsmanager dynamic reference to use the value stored in the secret to create the DB instance.
- D. Use the AWS::SecretsManager::Secret resource to generate a secure string. Store the secure string as a secret in AWS Secrets Manager. Use the secretsmanager dynamic reference to use the value stored in the secret to create the DB instance.

Correct Answer: D

Community vote distribution

D (78%)

B (22%)

MrTee **Highly Voted** 1 year, 9 months ago

Its a difficult choice between B and D

Option B leverages the existing AWS CLI command to generate a secure string, and then passes it as a parameter to CloudFormation, where it can be used to create the DB instance. But, if the use of Secrets Manager is already part of the organization's infrastructure, and the setup has already been completed, then option D may indeed be the simplest solution.

upvoted 7 times

sumanshu **Most Recent** 1 month ago

Selected Answer: D

A) Eliminated - This approach requires writing and maintaining custom Lambda code, which adds development effort.

B) Eliminated - This requires some setup in CodeBuild and manual handling of parameters, which adds complexity.

C) Eliminated - Same as option A (Reason)

D) Correct - Secrets Manager handles password generation and rotation automatically. This approach uses native CloudFormation functionality with no custom code.

upvoted 1 times

Saudis 4 months, 2 weeks ago

Selected Answer: D

Ans is B because the keyword is automatically generate passwords by secret manger by lest effort

upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

maurice2005 10 months, 3 weeks ago

Selected Answer: B

Where is the automatic generating of the password in option D?

upvoted 2 times

SerialiDr 1 year ago

Selected Answer: D

D. Use the AWS::SecretsManager::Secret resource to generate a secure string. Store the secure string as a secret in AWS Secrets Manager. Use the secretsmanager dynamic reference to use the value stored in the secret to create the DB instance: This solution efficiently uses AWS CloudFormation's native integration with AWS Secrets Manager. The AWS::SecretsManager::Secret resource type in CloudFormation can generate a secure string and store it as a secret. The secret value can then be used directly in the CloudFormation template to set the RDS instance password, using the secretsmanager dynamic reference. This approach minimizes development effort and leverages existing AWS services.

upvoted 3 times

fagilom 1 year, 1 month ago

D: This option leverages a native CloudFormation resource specifically designed for secret management. It eliminates the need for custom code or external tools, making it the simplest and most effort-efficient solution.

This approach minimizes custom code and utilizes native CloudFormation features, reducing overall complexity and maintenance.

upvoted 1 times

chewasa 1 year, 1 month ago

Selected Answer: D

you can create secrets with AWS::SecretsManager::Secret so it is the correct answer.

upvoted 2 times

LR2023 1 year, 1 month ago

Selected Answer: D

I was dilly dallying between B and D...but this helped me solidify my answer choice

https://docs.aws.amazon.com/secretsmanager/latest/userguide/cfn-example_reference-secret.html

upvoted 1 times

dezoito 1 year, 3 months ago

Selected Answer: D

With AWS CloudFormation, you can retrieve a secret to use in another AWS CloudFormation resource. A common scenario is to first create a secret with a password generated by Secrets Manager, and then retrieve the username and password from the secret to use as credentials for a new database.

https://docs.aws.amazon.com/secretsmanager/latest/userguide/cfn-example_reference-secret.html

upvoted 2 times

love777 1 year, 5 months ago

Selected Answer: B

Option B provides a straightforward approach to generating a secure string for the DB instance password and using it in CloudFormation with minimal development effort. Here's why this option is efficient:

CodeBuild Action: Using the AWS CodeBuild action within CodePipeline to generate a secure string using the `aws secretsmanager get-random-password` command allows you to easily create a random password without writing custom Lambda code.

CloudFormation Parameter: You can pass the generated secure string as a CloudFormation parameter with the `NoEcho` attribute set to `true`. This ensures that the parameter value won't be exposed in CloudFormation outputs or logs.

upvoted 4 times

FunkyFresco 1 year, 8 months ago

Selected Answer: D

The correct option is D. Create the password from secrets manager.

upvoted 4 times

delak 1 year, 8 months ago

Selected Answer: D

yes it's D

upvoted 2 times



rind2000 1 year, 8 months ago

Selected Answer: D

The answer is D

This is a secretsmanager dynamic reference sample in cloud formation

upvoted 2 times

  **chumji** 1 year, 8 months ago

I think answer is D

<https://aws.amazon.com/about-aws/whats-new/2022/12/amazon-rds-integration-aws-secrets-manager/>

upvoted 2 times

An organization is storing large files in Amazon S3, and is writing a web application to display meta-data about the files to end-users. Based on the metadata a user selects an object to download. The organization needs a mechanism to index the files and provide single-digit millisecond latency retrieval for the metadata.

What AWS service should be used to accomplish this?

- A. Amazon DynamoDB
- B. Amazon EC2
- C. AWS Lambda
- D. Amazon RDS

Correct Answer: A

Community vote distribution

A (100%)

MrTee **Highly Voted** 1 year, 9 months ago

Selected Answer: A

In this scenario, the metadata about the files can be stored in a DynamoDB table with a primary key based on the metadata attributes. This would enable the organization to quickly query and retrieve metadata about the files in real-time, with single-digit millisecond latency.

upvoted 14 times

sumanshu **Most Recent** 1 month ago

Selected Answer: A

DynamoDB is a fully managed NoSQL database service that provides single-digit millisecond latency for read and write operations.

D) Eliminated - While RDS can store metadata and index it using relational queries, it does not guarantee single-digit millisecond latency like DynamoDB.

upvoted 1 times

Saudis 3 months ago

Selected Answer: A

Keyword => single-digit millisecond latency retrieval for the metadata

upvoted 1 times

tsangckl 7 months, 2 weeks ago

This appear at 17 Jun exam

upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

SerialiDr 1 year ago

Selected Answer: A

A. Amazon DynamoDB: DynamoDB is a fast and flexible NoSQL database service that provides consistent single-digit millisecond latency for data retrieval. It is well-suited for applications that require high-performance data retrieval. The metadata of the files stored in S3 can be indexed and stored in a DynamoDB table, enabling efficient and quick access for the web application. This setup allows users to quickly browse metadata and select files for download.

upvoted 2 times

loctong 1 year, 8 months ago

Selected Answer: A

Amazon DynamoDB is a highly scalable and fully managed NoSQL database service that can provide fast and consistent performance at any scale. It is a suitable choice for indexing and storing metadata associated with files.

upvoted 3 times

A developer is creating an AWS Serverless Application Model (AWS SAM) template. The AWS SAM template contains the definition of multiple AWS Lambda functions, an Amazon S3 bucket, and an Amazon CloudFront distribution. One of the Lambda functions runs on Lambda@Edge in the CloudFront distribution. The S3 bucket is configured as an origin for the CloudFront distribution.

When the developer deploys the AWS SAM template in the eu-west-1 Region, the creation of the stack fails.

Which of the following could be the reason for this issue?

- A. CloudFront distributions can be created only in the us-east-1 Region.
- B. Lambda@Edge functions can be created only in the us-east-1 Region.
- C. A single AWS SAM template cannot contain multiple Lambda functions.
- D. The CloudFront distribution and the S3 bucket cannot be created in the same Region.

Correct Answer: B

Community vote distribution

B (93%)

7%

MrTee **Highly Voted** 1 year, 3 months ago

Selected Answer: B

it must be deployed to a region where Lambda@Edge is supported, such as us-east-1.
upvoted 12 times

zodraz **Highly Voted** 1 year, 2 months ago

Selected Answer: B

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/edge-functions-restrictions.html>

The Lambda function must be in the US East (N. Virginia) Region.
upvoted 9 times

sumanshu **Most Recent** 1 month ago

Selected Answer: B

A) Eliminated - CloudFront is a global service that is not tied to a specific region.
When creating a CloudFront distribution, the deployment can originate from any AWS region (e.g., eu-west-1).

B) Correct - Lambda@Edge functions must be deployed in the us-east-1 Region because they are globally replicated and integrated with CloudFront.
upvoted 1 times

65703c1 2 months, 1 week ago

Selected Answer: B

B is the correct answer.
upvoted 1 times

SD_CS 6 months ago

Selected Answer: B

B is the only answer that makes sense
upvoted 1 times

KarBiswa 7 months, 2 weeks ago

Selected Answer: B



<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-edge-how-it-works-tutorial.html>
clear mention
upvoted 1 times

tinyflame 12 months ago

Selected Answer: B

SAM can only specify one region
Langda@Edge only in us-east1 region

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-edge-how-it-works-tutorial.html>
upvoted 1 times

  **loctong** 1 year, 2 months ago

Selected Answer: C

Option A states that CloudFront distributions can only be created in the us-east-1 Region. This statement is incorrect because CloudFront distributions can be created in various AWS regions, including the eu-west-1 Region.
upvoted 2 times


A developer is integrating Amazon ElastiCache in an application. The cache will store data from a database. The cached data must populate real-time dashboards.

Which caching strategy will meet these requirements?

- A. A read-through cache
- B. A write-behind cache
- C. A lazy-loading cache
- D. A write-through cache

Correct Answer: D

Community vote distribution



MrTee Highly Voted 1 year, 9 months ago

Selected Answer: D

The best caching strategy for populating real-time dashboards using Amazon ElastiCache would be a write-through caching strategy. In this strategy, when new data is written to the database, it is also written to the cache. This ensures that the most current data is always available in the cache for the real-time dashboards to access, reducing the latency of the data retrieval. Additionally, using a write-through cache ensures that data consistency is maintained between the database and the cache, as any changes to the data are written to both locations simultaneously.

upvoted 16 times

sumanshu Most Recent 1 month ago

Selected Answer: D

A) Eliminated

In a read-through cache, when an application tries to retrieve data:

If the data is not available in the cache, the cache retrieves it from the database, stores it in the cache, and returns it to the application.

If the data is available in the cache, it is returned directly.

The cache is updated only when data is requested, which could result in outdated data being served to real-time dashboards if no recent read request has occurred.

upvoted 1 times

sumanshu 1 month ago

B) Eliminated

In a write-behind cache, data is written to the cache first, and the database is updated asynchronously (in the background).

This caching strategy prioritizes write performance but risks data loss in the event of a failure because the cache may not yet have updated the database.

For real-time dashboards, there is no guarantee that the cache is always up-to-date because the data is updated asynchronously.

upvoted 1 times

sumanshu 1 month ago

C) Eliminated

In a lazy-loading cache, the data is loaded into the cache only when requested.

If the data is not in the cache (a "cache miss"), it is fetched from the database, loaded into the cache, and returned to the application.

While this reduces unnecessary caching, it does not proactively update the cache. Therefore, data in the cache may become stale, which is unsuitable for real-time dashboards.

upvoted 1 times

sumanshu 1 month ago

D) Correct

In a write-through cache, every write operation to the database is immediately mirrored in the cache.

This ensures that the cache always contains the most up-to-date data and is ideal for scenarios where real-time consistency is required, such as updating dashboards.

upvoted 1 times

🗨️ 👤 **Saudis** 3 months ago

Selected Answer: D

the keyword => real-time

upvoted 1 times

🗨️ 👤 **tirthyakamaldasgupta** 5 months, 3 weeks ago

When using a write-through cache strategy, the cache is updated in real-time alongside the database. This ensures that the cached data remains consistent with the underlying database. According to AWS best practices, this approach pushes data into the cache at the time it is written to the database, reducing the risk of serving stale data.

In contrast, option A (a read-through cache strategy) may result in stale data, particularly if the cache has a time-to-live (TTL) setting that allows data to remain in the cache longer than it remains accurate in the database. This can be problematic for real-time dashboards that require up-to-date information. For these reasons, I opted for option D.

<https://aws.amazon.com/caching/best-practices/#:~:text=Write%2Dthrough,also%20pushed%20into%20the%20cache.>

Please correct me if my understanding is incorrect, as I am still learning.

upvoted 1 times

🗨️ 👤 **65703c1** 8 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

🗨️ 👤 **Walker17** 12 months ago

B. Write Behind Cache.

upvoted 1 times

🗨️ 👤 **SerialiDr** 1 year ago

Selected Answer: D

D. A write-through cache: A write-through caching strategy immediately writes data to both the cache and the database at the same time. This approach ensures that the cache always contains the most recent data, making it highly suitable for applications that require up-to-date information, such as real-time dashboards.

upvoted 2 times

🗨️ 👤 **tqiu654** 1 year, 1 month ago

Selected Answer: C

ChatGPT:C

upvoted 1 times

🗨️ 👤 **Prem28** 1 year, 7 months ago

ans- A

Option D, a write-through cache, is incorrect because it would not meet the requirement of populating real-time dashboards. A write-through cache writes data to the cache and the database at the same time. This means that the data in the cache would always be up-to-date, but it would also mean that the cache would always be lagging behind the database. This would cause a delay in populating real-time dashboards.

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

I agree. I think it's A because D is better option when you need data to be consistent and highly available since data is always up to date but as Prem28 says it lags behind on latency when compared to read-through. What I get from the question is they need strategy for "real-time" dashboards --> reduction of latency not accuracy or consistent data

upvoted 1 times

🗨️ 👤 **loctong** 1 year, 8 months ago

Selected Answer: D

A write-through cache strategy involves writing data to both the cache and the underlying database simultaneously. When data is updated or inserted into the database, it is also stored or updated in the cache to ensure that the cache remains up-to-date with the latest data.

upvoted 2 times

A developer is creating an AWS Lambda function. The Lambda function needs an external library to connect to a third-party solution. The external library is a collection of files with a total size of 100 MB. The developer needs to make the external library available to the Lambda execution environment and reduce the Lambda package space.

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

- A. Create a Lambda layer to store the external library. Configure the Lambda function to use the layer.
- B. Create an Amazon S3 bucket. Upload the external library into the S3 bucket. Mount the S3 bucket folder in the Lambda function. Import the library by using the proper folder in the mount point.
- C. Load the external library to the Lambda function's /tmp directory during deployment of the Lambda package. Import the library from the /tmp directory.
- D. Create an Amazon Elastic File System (Amazon EFS) volume. Upload the external library to the EFS volume. Mount the EFS volume in the Lambda function. Import the library by using the proper folder in the mount point.

Correct Answer: A

Community vote distribution

A (100%)

 **MrTee** Highly Voted 1 year, 9 months ago

Selected Answer: A

Create a Lambda layer to store the external library. Configure the Lambda function to use the layer. This will allow the developer to make the external library available to the Lambda execution environment without having to include it in the Lambda package, which will reduce the Lambda package space. Using a Lambda layer is a simple and straightforward solution that requires minimal operational overhead.

upvoted 13 times

 **sumanshu** Most Recent 1 month ago


Selected Answer: A

Layers are managed separately from the Lambda function.

They reduce the size of the deployment package.

Layers are easy to use, and there is no need for complex infrastructure or manual management.


upvoted 1 times

 **Saurabh04** 5 months, 3 weeks ago

Option B is correct because it is straightforward with lesser operation overhead than managing layers. Option A and C are incorrect. While Option A approach allows you to separate library from your function code, it introduces some operational overhead in managing layers.


Option C is simple but doesn't separate library from your function code.

upvoted 1 times

 **tsangckl** 7 months, 2 weeks ago


This appear at 17 Jun exam

upvoted 1 times

 **frangesk** 5 months, 1 week ago

Do you remember the answer?

upvoted 1 times

 **65703c1** 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

 **KillThemWithKindness** 11 months, 2 weeks ago

Selected Answer: A

You can add up to five layers to a Lambda function. The total unzipped size of the function and all layers cannot exceed the unzipped deployment package size quota of 250 MB. For more information, see Lambda quotas.

upvoted 4 times

SerialIDr 1 year ago

Selected Answer: A

A. Create a Lambda layer to store the external library. Configure the Lambda function to use the layer: This is the most suitable solution. Lambda layers allow you to include libraries and other dependencies without including them in the deployment package of your Lambda function. By creating a layer with the external library and configuring the Lambda function to use this layer, the developer can easily manage and update the library independently of the Lambda function code, reducing the package size and operational overhead.

upvoted 2 times

CalvinL4 1 year ago

One lambda layer only allows 50 mb for storage. The file is 100 MB. So I will vote for D unless the library can break down into less than 5 layers.

upvoted 1 times

loctong 1 year, 8 months ago

Selected Answer: A

By creating a Lambda layer, you can separate the external library from the Lambda function code itself and make it available to multiple functions. This approach offers the following benefits:

upvoted 2 times

dan80 1 year, 9 months ago

Selected Answer: A

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-layers.html>

upvoted 3 times

A company has a front-end application that runs on four Amazon EC2 instances behind an Elastic Load Balancer (ELB) in a production environment that is provisioned by AWS Elastic Beanstalk. A developer needs to deploy and test new application code while updating the Elastic Beanstalk platform from the current version to a newer version of Node.js. The solution must result in zero downtime for the application.

Which solution meets these requirements?

- A. Clone the production environment to a different platform version. Deploy the new application code, and test it. Swap the environment URLs upon verification.
- B. Deploy the new application code in an all-at-once deployment to the existing EC2 instances. Test the code. Redeploy the previous code if verification fails.
- C. Perform an immutable update to deploy the new application code to new EC2 instances. Serve traffic to the new instances after they pass health checks.
- D. Use a rolling deployment for the new application code. Apply the code to a subset of EC2 instances until the tests pass. Redeploy the previous code if the tests fail.

Correct Answer: C

Community vote distribution



MrTee **Highly Voted** 1 year, 9 months ago

Selected Answer: C

Option C is the correct solution that meets the requirements. Performing an immutable update to deploy the new application code to new EC2 instances and serving traffic to the new instances after they pass health checks will ensure zero downtime for the application.

Option A would work but cloning the production environment to a different platform version will result in a longer deployment time and can impact the cost of the environment.

upvoted 23 times

awsdummie 1 year, 9 months ago

C is incorrect, after passing health checks the elastic Beanstalk transfers them to the original Auto Scaling group. No testing or platform update is done.

upvoted 5 times

yeacuz 1 year, 8 months ago

I would agree that option A can affect the cost, but cost is not the issue. The question is asking for zero downtime. I believe the answer is option A

upvoted 3 times

gagol14 **Highly Voted** 1 year, 7 months ago

Selected Answer: A

Not C: While an immutable update can ensure zero downtime during the deployment process, it doesn't account for updating the Elastic Beanstalk platform version.

upvoted 9 times

sumanshu **Most Recent** 1 month ago

Selected Answer: C

A) Eliminated - This solution meets the requirement of zero downtime but may have additional resource and cost overhead.

B) Eliminated - All-at-once deployment means that all instances are updated at the same time, which would result in downtime during the deployment.

C) Correct - Immutable updates ensure that new EC2 instances are created, and only healthy instances will serve traffic. The old EC2 instances are not affected until the new instances are confirmed to be healthy and live.

D) Eliminated - Rolling deployments help ensure that some EC2 instances are always running the old code and serving traffic while others are being updated.

upvoted 1 times

🗨️ 👤 **MasoudK** 4 months ago

Option A is Correct not C: By cloning the production environment to a different platform version, you create a separate environment where you can safely deploy and test the new application code and platform version without affecting the live production environment. • Option C (Immutable update): While immutable updates ensure zero downtime by deploying to new instances, they do not address the need to update the Elastic Beanstalk platform version. Additionally, this approach can be more resource-intensive and costly.

upvoted 1 times

🗨️ 👤 **Saudis** 4 months, 2 weeks ago

Selected Answer: C

Ans is c

upvoted 1 times

🗨️ 👤 **65703c1** 8 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗨️ 👤 **ibratoev** 10 months, 1 week ago

It is A: https://docs.amazonaws.cn/en_us/elasticbeanstalk/latest/dg/using-features.platform.upgrade.html#using-features.platform.upgrade.config

upvoted 2 times

🗨️ 👤 **KarBiswa** 10 months, 2 weeks ago

This question must be true for 2 options because C & D are both correct

upvoted 1 times

🗨️ 👤 **KarBiswa** 10 months, 3 weeks ago

Selected Answer: D

https://docs.amazonaws.cn/en_us/elasticbeanstalk/latest/dg/using-features.rolling-version-deploy.html

upvoted 1 times

🗨️ 👤 **SerialiDr** 1 year ago

Selected Answer: C

The solutions that best meet the requirements for zero downtime are:

A. Clone the production environment to a different platform version. Deploy the new application code, and test it. Swap the environment URLs upon verification.

C. Perform an immutable update to deploy the new application code to new EC2 instances. Serve traffic to the new instances after they pass health checks.

Both options A and C provide robust strategies for deploying updates with zero downtime, allowing for thorough testing in an isolated environment before directing production traffic to the new setup.

upvoted 3 times

🗨️ 👤 **Certified101** 1 year, 1 month ago

Selected Answer: A

Not C: It doesn't account for updating the Elastic Beanstalk platform version. This would affect both the live and test environments.

It's also best practice to have 2 separate environments for production and test and there is no mention of cost optimisation here.

upvoted 3 times

🗨️ 👤 **tqiu654** 1 year, 1 month ago

Selected Answer: A

ChatGPT: A

upvoted 2 times

🗨️ 👤 **Rameez1** 1 year, 3 months ago

Selected Answer: C

A & C both work for given scenario but C does it more feasibly for Elastic Beanstalk with zero downtime.

upvoted 1 times

🗨️ 👤 **stilloneway** 1 year, 5 months ago

Selected Answer: C

Key terminology in question is "Test". So it should be immutable for quick rollback in case of test not working.

upvoted 2 times

  **[Removed]** 1 year, 1 month ago



Option A offers quick rollback too... did some research and cloning is same as blue/green deployments. with that said, I think the answer is A

upvoted 1 times

  **CrescentShared** 1 year ago

It's a downtime if test fails and rollback.

upvoted 1 times

  **love777** 1 year, 5 months ago



Selected Answer: C

Explanation:

Immutable Update with Elastic Beanstalk:

With an immutable update, Elastic Beanstalk provisions new instances with the updated code while keeping the existing instances running. The traffic is shifted gradually to the new instances after they pass health checks, ensuring that there is no downtime during the deployment. If any issue arises during the deployment, traffic is still being served by the existing instances.

upvoted 4 times

  **Naj_64** 1 year, 5 months ago



Selected Answer: D

Screenshot of Step 4 of Method 1 in the link:

https://docs.amazonaws.cn/en_us/elasticbeanstalk/latest/dg/using-features.platform.upgrade.html#using-features.platform.upgrade.config

"...your application is unavailable during the update. To keep at least one instance in service during the update, enable rolling updates"

upvoted 2 times

  **Naj_64** 1 year, 5 months ago

I take this back. I'm going with A

"However, you can avoid this downtime by deploying the new version to a separate environment. The existing environment's configuration is copied and used to launch the green environment with the new version of the application. The new green environment will have its own URL. When it's time to promote the green environment to serve production traffic, you can use Elastic Beanstalk's Swap Environment URLs feature."

<https://docs.aws.amazon.com/whitepapers/latest/blue-green-deployments/swap-the-environment-of-an-elastic-beanstalk-application.html>

upvoted 1 times

  **MG1407** 1 year, 5 months ago

Selected Answer: A

A is the answer. Sorry about the double post ...

https://docs.amazonaws.cn/en_us/elasticbeanstalk/latest/dg/using-features.platform.upgrade.html#using-features.platform.upgrade.config

upvoted 4 times

A developer is creating an AWS Lambda function. The Lambda function will consume messages from an Amazon Simple Queue Service (Amazon SQS) queue. The developer wants to integrate unit testing as part of the function's continuous integration and continuous delivery (CI/CD) process.

How can the developer unit test the function?

- A. Create an AWS CloudFormation template that creates an SQS queue and deploys the Lambda function. Create a stack from the template during the CI/CD process. Invoke the deployed function. Verify the output.
- B. Create an SQS event for tests. Use a test that consumes messages from the SQS queue during the function's CI/CD process.
- C. Create an SQS queue for tests. Use this SQS queue in the application's unit test. Run the unit tests during the CI/CD process.
- D. Use the `aws lambda invoke` command with a test event during the CI/CD process.

Correct Answer: D

Community vote distribution

D (42%) C (38%) B (20%)

 **gagol14** Highly Voted 1 year, 7 months ago

Selected Answer: C


Unit testing is a type of testing that verifies the correctness of individual units of source code, typically functions or methods. When unit testing a Lambda function that interacts with Amazon SQS, you can create a separate test SQS queue that the Lambda function interacts with during testing. You would then validate the behavior of the function based on its interactions with the test queue. This approach isolates the function's behavior from the rest of the system, which is a key principle of unit testing.

Option A is incorrect because AWS CloudFormation is typically used for infrastructure deployment, not for unit testing.

Option B is incorrect because it does not actually test the function; it only creates an event.

Option D is incorrect because the `'aws lambda invoke'` command is used to manually trigger a Lambda function, but doesn't necessarily facilitate testing the function's behavior when consuming messages from an SQS queue.

upvoted 16 times

 **redfivedog** Highly Voted 1 year, 6 months ago

Selected Answer: D

D is correct here. Both B and C are integration tests as they are using an actual SQS queue in the tests and not mocking it out.

upvoted 13 times

 **mooncake1** Most Recent 6 days, 11 hours ago

Selected Answer: D

Those who select C and still think they are right - Should study more about unit test & integrate testing before AWS

upvoted 1 times

 **sumanshu** 1 month ago

Selected Answer: D

A) Eliminated - This approach involves creating and destroying AWS resources (SQS, Lambda) for each test, which increases the time, complexity, and potential costs associated with running tests.

B) Eliminated - t involves real AWS resources (SQS)

C) Eliminated - Even though it's meant for testing, you're still using an actual AWS service (SQS) to conduct the test, making it less of a unit test and more of an integration test.

D) Correct - You use the `aws lambda invoke` command to invoke the Lambda function directly during the test process, without needing to rely on actual resources like an SQS queue.

upvoted 1 times

🗨️ **ShakthiGCP** 1 month, 2 weeks ago

Selected Answer: C

Unit testing is a type of testing that verifies the correctness of individual units of source code, typically functions or methods

upvoted 1 times

🗨️ **f271c23** 1 month, 3 weeks ago

Selected Answer: C

option C seems more appropriate. As there is a clear isolation of test stage environment and it is staged prior to the deployment.

Option D triggered manually and consuming from the same queue that was meant for the Production environment

upvoted 1 times

🗨️ **CloudChingon** 2 months, 2 weeks ago

Selected Answer: D

It's testing so option D seems more logic. The other options would put a message in SQS.

upvoted 1 times

🗨️ **Anandesh** 6 months, 1 week ago

Selected Answer: D

<https://docs.aws.amazon.com/lambda/latest/dg/testing-guide.html>

upvoted 1 times

🗨️ **65703c1** 8 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

🗨️ **Morales** 9 months, 3 weeks ago

Selected Answer: D

In the case of unit tests, whose objective is to isolate the tested unit, option D is the one that most isolates the unit.

upvoted 1 times

🗨️ **41eb566** 10 months, 2 weeks ago

Selected Answer: C

To unit test the AWS Lambda function that consumes messages from an Amazon SQS queue as part of the CI/CD process, the developer can follow option C:

C. Create an SQS queue for tests. Use this SQS queue in the application's unit test. Run the unit tests during the CI/CD process.

upvoted 1 times

🗨️ **KillThemWithKindness** 11 months, 2 weeks ago

Selected Answer: D

In production, our Lambda function code will directly access the AWS resources we defined in our function handler; however, in our unit tests we want to isolate our code and replace the AWS resources with simulations. This isolation facilitates running unit tests in an isolated environment to prevent accidental access to actual cloud resources.

<https://aws.amazon.com/blogs/devops/unit-testing-aws-lambda-with-python-and-mock-aws-services/>

upvoted 1 times

🗨️ **SerialiDr** 1 year ago

Selected Answer: D

D. Use the aws lambda invoke command with a test event during the CI/CD process: This option is closer to what unit testing entails. The aws lambda invoke command can be used to invoke the Lambda function with a simulated event payload that mimics an SQS message. This allows the developer to test the function's logic and handling of SQS messages without needing an actual SQS queue. The test can focus on how the function processes the input and generates output, which is the essence of unit testing.

upvoted 3 times

🗨️ **CrescentShared** 1 year ago

Anybody find this question in the exam, please? The question itself looks so wrong to me, the action of testing the lambda function does not seem like a 'unit test' already... Isn't the unit test testing all the Classes inside the lambda function?

upvoted 2 times

🗨️ **Certified101** 1 year, 1 month ago

Selected Answer: C

C there should be a separate isolated test environment

D will only invoke the lambda and not test SQS polling.

upvoted 2 times

🗨️ 👤 **tqiu654** 1 year, 1 month ago

Selected Answer: D

ChatGPT:D

upvoted 1 times

🗨️ 👤 **ShawnWon** 1 year, 2 months ago

B.

Option A (CloudFormation template for SQS queue and Lambda function) involves more of an integration test rather than a unit test. It's typically preferable to keep unit tests isolated and focused on the specific functionality of the function.

Option C (Create an SQS queue for tests) might involve additional setup and cleanup steps, and it could introduce dependencies that impact the isolation of unit tests.

Option D (aws lambda invoke command with a test event) is similar to Option B, but creating a test event is generally more flexible and allows for a clearer representation of the expected input to the Lambda function.

upvoted 3 times

A developer is working on a web application that uses Amazon DynamoDB as its data store. The application has two DynamoDB tables: one table that is named artists and one table that is named songs. The artists table has artistName as the partition key. The songs table has songName as the partition key and artistName as the sort key.

The table usage patterns include the retrieval of multiple songs and artists in a single database operation from the webpage. The developer needs a way to retrieve this information with minimal network traffic and optimal application performance.

Which solution will meet these requirements?

- A. Perform a BatchGetItem operation that returns items from the two tables. Use the list of songName/artistName keys for the songs table and the list of artistName key for the artists table.
- B. Create a local secondary index (LSI) on the songs table that uses artistName as the partition key. Perform a query operation for each artistName on the songs table that filters by the list of songName. Perform a query operation for each artistName on the artists table.
- C. Perform a BatchGetItem operation on the songs table that uses the songName/artistName keys. Perform a BatchGetItem operation on the artists table that uses artistName as the key.
- D. Perform a Scan operation on each table that filters by the list of songName/artistName for the songs table and the list of artistName in the artists table.

Correct Answer: A

Community vote distribution

A (90%)

5%

 **csG13** Highly Voted 1 year, 7 months ago

Selected Answer: A

The correct answer is A. BatchGetItem can return one or multiple items from one or more tables. For reference check the link below

https://docs.aws.amazon.com/amazondynamodb/latest/APIReference/API_BatchGetItem.html

upvoted 8 times

 **sumanshu** Most Recent 1 month ago

Selected Answer: A

BatchGetItem allows fetching multiple items from one or more DynamoDB tables in a single call, which reduces network traffic.

C) Eliminated - While each operation is efficient within its table, two separate BatchGetItem calls still result in more round trips to DynamoDB than a single batch operation for both tables

D) Eliminated - Scan operations are very inefficient because they read every item in the table, even if only a subset of data is needed.

upvoted 1 times

 **65703c1** 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

 **KarBiswa** 11 months, 1 week ago

Selected Answer: C

I would go for it because typically we are taking the advantage of key selection

upvoted 1 times

 **NSA_Poker** 4 months, 2 weeks ago

(C) performs BatchGetItem operation twice; that's NOT optimal and above the minimal network traffic necessary.

Amazon DynamoDB BatchGetItem - returns the attributes of one or more items from one or more tables. You identify requested items by primary key.

upvoted 1 times

 **KarBiswa** 11 months, 1 week ago

Sorry its Option A saying multiple songs so list will be right option

upvoted 1 times

🗨️ **SerialiDr** 1 year ago

Selected Answer: A

The BatchGetItem API allows you to get up to 100 items from one or more DynamoDB tables in a single operation, which can reduce the number of network requests. This is efficient for retrieving a specific list of items when you know the primary keys (partition key and sort key, if applicable) of the items you want to retrieve.

upvoted 4 times

🗨️ **norris81** 1 year, 4 months ago

Selected Answer: A

https://docs.aws.amazon.com/amazondynamodb/latest/APIReference/API_BatchGetItem.html

upvoted 2 times

🗨️ **rind2000** 1 year, 6 months ago

Selected Answer: B

Agree 100% with Caiyi.

upvoted 1 times

🗨️ **caiyi** 1 year, 7 months ago

B.

By creating a local secondary index (LSI) on the songs table with artistName as the partition key, you can efficiently query the songs table for each artistName in the list of artists. This approach allows you to retrieve the desired songs for multiple artists with minimal network traffic.

upvoted 3 times

🗨️ **GripZA** 1 year, 5 months ago

You can't create a LSI on an existing DDB table

upvoted 6 times

🗨️ **remynick** 1 year, 5 months ago

I dont agree, we need to creat a global secondary index to use artistName as the partition ke

upvoted 2 times

🗨️ **Baba_Eni** 1 year, 7 months ago

Selected Answer: A

https://docs.aws.amazon.com/amazondynamodb/latest/APIReference/API_BatchGetItem.html

upvoted 3 times

A company is developing an ecommerce application that uses Amazon API Gateway APIs. The application uses AWS Lambda as a backend. The company needs to test the code in a dedicated, monitored test environment before the company releases the code to the production environment.

Which solution will meet these requirements?

- A. Use a single stage in API Gateway. Create a Lambda function for each environment. Configure API clients to send a query parameter that indicates the environment and the specific Lambda function.
- B. Use multiple stages in API Gateway. Create a single Lambda function for all environments. Add different code blocks for different environments in the Lambda function based on Lambda environment variables.
- C. Use multiple stages in API Gateway. Create a Lambda function for each environment. Configure API Gateway stage variables to route traffic to the Lambda function in different environments.
- D. Use a single stage in API Gateway. Configure API clients to send a query parameter that indicates the environment. Add different code blocks for different environments in the Lambda function to match the value of the query parameter.

Correct Answer: C

Community vote distribution

C (100%)

 **csG13** Highly Voted 1 year, 7 months ago

Selected Answer: C

The answer is C - we should create multiple stages and different Lambdas that will be utilised based on API Gateway stages variables.


<https://docs.aws.amazon.com/apigateway/latest/developerguide/amazon-api-gateway-using-stage-variables.html>

upvoted 14 times

 **tsangckl** Most Recent 7 months, 2 weeks ago

This appear at 17 Jun exam

upvoted 1 times

 **65703c1** 8 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

 **SerialiDr** 1 year ago

Selected Answer: C

C. Use multiple stages in API Gateway. Create a Lambda function for each environment. Configure API Gateway stage variables to route traffic to the Lambda function in different environments: This is the recommended approach. Using multiple stages in API Gateway (one for testing and one for production) allows for clear separation of environments. Having a dedicated Lambda function for each environment ensures isolation and reduces the risk of accidental changes impacting the production environment. API Gateway stage variables can be used to manage configurations specific to each stage, such as function names or other parameters.

upvoted 3 times

A developer creates an AWS Lambda function that retrieves and groups data from several public API endpoints. The Lambda function has been updated and configured to connect to the private subnet of a VPC. An internet gateway is attached to the VPC. The VPC uses the default network ACL and security group configurations.

The developer finds that the Lambda function can no longer access the public API. The developer has ensured that the public API is accessible, but the Lambda function cannot connect to the API

How should the developer fix the connection issue?

- A. Ensure that the network ACL allows outbound traffic to the public internet.
- B. Ensure that the security group allows outbound traffic to the public internet.
- C. Ensure that outbound traffic from the private subnet is routed to a public NAT gateway.
- D. Ensure that outbound traffic from the private subnet is routed to a new internet gateway.

Correct Answer: C

Community vote distribution

C (100%)

 **Dushank** Highly Voted 10 months, 3 weeks ago

Selected Answer: C

When a Lambda function is configured to connect to a VPC, it loses its default internet access. To allow the Lambda function to access the public internet, it must be connected to a private subnet in the VPC that is configured to route its traffic through a NAT Gateway (Network Address Translation Gateway).

The Internet Gateway is usually used to provide internet access to resources in the public subnet, but for resources in the private subnet, a NAT Gateway is required.


upvoted 6 times

 **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

 **SerialiDr** 6 months, 2 weeks ago

Selected Answer: C

C. Ensure that outbound traffic from the private subnet is routed to a public NAT gateway: This is the most likely solution. Lambda functions in a private subnet require a NAT (Network Address Translation) gateway or NAT instance in a public subnet to access the public internet, as private subnets do not have direct internet access. The VPC route table associated with the private subnet needs to have a route that directs internet-bound traffic to the NAT gateway.

upvoted 2 times

 **Naj_64** 11 months, 2 weeks ago

Selected Answer: C

NAT Gateway from a public subnet is required.

upvoted 1 times

 **cmonthatsme** 12 months ago

Selected Answer: C

The Lambda function is running in a private subnet of the VPC, it needs to send outbound traffic to the internet to reach the API endpoints. To enable this, a NAT gateway is required.

upvoted 1 times

 **Parsons** 12 months ago

Selected Answer: C

C is correct.

with Lambda, You need an IP of NAT GW to be able to access public internet.

upvoted 1 times

  **cloudenthusiast** 12 months ago

Selected Answer: C

it leverages a NAT gateway, which is a service that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating a connection with those instances.

upvoted 2 times

A developer needs to store configuration variables for an application. The developer needs to set an expiration date and time for the configuration. The developer wants to receive notifications before the configuration expires.

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

- A. Create a standard parameter in AWS Systems Manager Parameter Store. Set Expiration and ExpirationNotification policy types.
- B. Create a standard parameter in AWS Systems Manager Parameter Store. Create an AWS Lambda function to expire the configuration and to send Amazon Simple Notification Service (Amazon SNS) notifications.
- C. Create an advanced parameter in AWS Systems Manager Parameter Store. Set Expiration and ExpirationNotification policy types.
- D. Create an advanced parameter in AWS Systems Manager Parameter Store. Create an Amazon EC2 instance with a cron job to expire the configuration and to send notifications.

Correct Answer: C

Community vote distribution

C (81%)

Other

Parsons Highly Voted 1 year, 5 months ago

Selected Answer: C

C is correct.

You have to use "advanced parameter in AWS Systems Manager Parameter Store" to be able to Set Expiration and ExpirationNotification policy types.

upvoted 11 times

Anandesh Most Recent 6 months, 2 weeks ago

Selected Answer: C

<https://docs.aws.amazon.com/systems-manager/latest/userguide/parameter-store-policies.html>

upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

KarBiswa 11 months, 1 week ago

Selected Answer: C

<https://docs.aws.amazon.com/systems-manager/latest/userguide/parameter-store-policies.html>

upvoted 1 times

Trung125122 11 months, 1 week ago

Selected Answer: A

A is sufficient. C is abundant

upvoted 1 times

NijeshT 1 year, 2 months ago

Advanced Parameters: These offer more capabilities, such as adding policies for expiration and triggering notifications

upvoted 1 times

Rameez1 1 year, 3 months ago

Selected Answer: B

Using Lambda function and SNS will address the requirement with least operational overhead.

upvoted 2 times

Rameez1 1 year, 3 months ago

Changing my mind option A is correct here.

upvoted 2 times

Fizbo 1 year, 1 month ago

It is C. standard tier does not have those features

upvoted 1 times

🗨️ 👤 **Gold07** 1 year, 3 months ago

A is the right Answer

upvoted 2 times

🗨️ 👤 **worseforwear** 1 year, 5 months ago

Selected Answer: C

You can't set expiration policy on standard parameter

upvoted 4 times

🗨️ 👤 **cmonthatsme** 1 year, 5 months ago

Selected Answer: A

By creating a standard parameter, you can set an expiration date for the parameter

upvoted 2 times

🗨️ 👤 **cloudenthusiast** 1 year, 6 months ago

Selected Answer: C

it leverages the advanced parameter tier and the parameter policies feature of Parameter Store, which meet the requirements with the least operational overhead.

upvoted 4 times

A company is developing a serverless application that consists of various AWS Lambda functions behind Amazon API Gateway APIs. A developer needs to automate the deployment of Lambda function code. The developer will deploy updated Lambda functions with AWS CodeDeploy. The deployment must minimize the exposure of potential errors to end users. When the application is in production, the application cannot experience downtime outside the specified maintenance window.

Which deployment configuration will meet these requirements with the LEAST deployment time?

- A. Use the AWS CodeDeploy in-place deployment configuration for the Lambda functions. Shift all traffic immediately after deployment.
- B. Use the AWS CodeDeploy linear deployment configuration to shift 10% of the traffic every minute.
- C. Use the AWS CodeDeploy all-at-once deployment configuration to shift all traffic to the updated versions immediately.
- D. Use the AWS CodeDeploy predefined canary deployment configuration to shift 10% of the traffic immediately and shift the remaining traffic after 5 minutes.

Correct Answer: D

Community vote distribution

D (90%)

10%

65703c1 2 months, 1 week ago

Selected Answer: D

D is the correct answer.
upvoted 2 times

KarBiswa 5 months, 1 week ago

Selected Answer: D

<https://docs.aws.amazon.com/whitepapers/latest/practicing-continuous-integration-continuous-delivery/deployment-methods.html#:~:text=A%20variation%20of,is%20gradually%20increased>.
upvoted 2 times

rimaSamir 6 months ago

Selected answer is A.

To them who have chosen D, you have forgotten also about "When the application is in production, the application cannot experience downtime outside the specified maintenance window."
upvoted 1 times

SerialiDr 6 months, 2 weeks ago

Selected Answer: D

D. Use the AWS CodeDeploy predefined canary deployment configuration to shift 10% of the traffic immediately and shift the remaining traffic after 5 minutes: The canary deployment strategy first shifts a small percentage of traffic to the new version (e.g., 10%) and, after a specified period (e.g., 5 minutes), shifts the remaining traffic. This approach allows for initial validation of the new version with minimal user exposure before full rollout, balancing speed and risk mitigation.
upvoted 2 times

c9ebec2 7 months, 2 weeks ago

Selected Answer: D

Lambda deploy supports just Linear or Canary. So answer is D. Linear or All
upvoted 2 times

aravindpti 8 months ago

Answer A.

<https://aws.amazon.com/blogs/containers/aws-codedeploy-now-supports-linear-and-canary-deployments-for-amazon-ecs/>
upvoted 1 times

jingle4944 9 months, 1 week ago



Canary deployment is supported: <https://aws.amazon.com/blogs/compute/implementing-safe-aws-lambda-deployments-with-aws-codedeploy/>
upvoted 1 times

passhojaun 9 months, 2 weeks ago

Selected Answer: A



Canary is not supported in AWS CodeDeploy.

upvoted 1 times

  **Jaimoo** 9 months, 1 week ago



<https://aws.amazon.com/es/blogs/containers/aws-codedeploy-now-supports-linear-and-canary-deployments-for-amazon-ecs/>

upvoted 2 times

  **passhojaun** 9 months, 2 weeks ago

Canary is not supported in AWS CodeDeploy.

upvoted 2 times

  **Monivs** 6 months, 4 weeks ago

Canary is supported by code deploy

<https://docs.aws.amazon.com/codedeploy/latest/userguide/welcome.html>



upvoted 1 times

  **Yuxing_Li** 11 months, 1 week ago

Selected Answer: D

Canary is faster than linear in this case.

upvoted 2 times

  **love777** 11 months, 1 week ago

Selected Answer: A


Explanation:

In an AWS Lambda context, using the in-place deployment configuration minimizes deployment time and provides fast updates to the function's code. In this case, the application consists of AWS Lambda functions behind Amazon API Gateway APIs. With the in-place deployment configuration, all traffic is shifted to the updated versions of the Lambda functions immediately after deployment.

Option B suggests a linear deployment configuration that shifts 10% of the traffic every minute. While this provides controlled deployment and gradual rollout, it might not be the fastest approach if you want to minimize deployment time.

Option C suggests an all-at-once deployment configuration. While this configuration might be fast, it poses a higher risk of exposing potential errors to end users all at once.

upvoted 1 times

  **Monivs** 6 months, 4 weeks ago

Inplace deployment is not supported by ECS and Lambda

upvoted 1 times

  **RaidenKurosaki** 12 months ago

Selected Answer: D

Canary deployment

upvoted 2 times

  **Parsons** 12 months ago

Selected Answer: D

D is correct.

Keyword:

- "must minimize the exposure of potential errors to end users", you just have to trade-off 10% of traffic

- "cannot experience downtime ", eliminate C.

- "LEAST deployment time", with B, You have to take 10 mins other than D just 5 min.

upvoted 4 times

  **cloudenthusiast** 12 months ago

Selected Answer: D

the predefined canary deployment configuration, which shifts a small percentage of traffic to the updated versions immediately, and then shifts the remaining traffic after a specified period

upvoted 2 times

A company created four AWS Lambda functions that connect to a relational database server that runs on an Amazon RDS instance. A security team requires the company to automatically change the database password every 30 days.

Which solution will meet these requirements MOST securely?

- A. Store the database credentials in the environment variables of the Lambda function. Deploy the Lambda function with the new credentials every 30 days.
- B. Store the database credentials in AWS Secrets Manager. Configure a 30-day rotation schedule for the credentials.
- C. Store the database credentials in AWS Systems Manager Parameter Store secure strings. Configure a 30-day schedule for the secure strings.
- D. Store the database credentials in an Amazon S3 bucket that uses server-side encryption with customer-provided encryption keys (SSE-C). Configure a 30-day key rotation schedule for the customer key.

Correct Answer: B

Community vote distribution

B (100%)

 **Dushank**  1 year, 4 months ago

Selected Answer: B


The most secure and automated way to handle database credential rotation is to use AWS Secrets Manager. Secrets Manager can automatically rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. You can configure Secrets Manager to automatically rotate the secrets for you according to a schedule you specify, making it easier to adhere to best practices for security.

upvoted 6 times

 **tsangcki**  7 months, 2 weeks ago

This appear at 17 Jun exam

upvoted 1 times

 **65703c1** 8 months, 1 week ago

Selected Answer: B

B is the correct answer.


upvoted 1 times

 **RaidenKurosaki** 1 year, 5 months ago

Selected Answer: B

Secrets Manager supports auto rotation. Systems Manager does not do that.

upvoted 4 times

 **Parsons** 1 year, 5 months ago

Selected Answer: B

B is correct.

Keyword: "automatically change the database password every 30 days"


upvoted 2 times

 **cloudenthusiast** 1 year, 6 months ago

Selected Answer: B

Secrets Manager supports automatic rotation of secrets by using either built-in or custom Lambda functions

upvoted 3 times

 **niks1221** 1 year, 5 months ago

Did you give your exam recently?

If yes, how many questions were from here?

upvoted 1 times

A developer is setting up a deployment pipeline. The pipeline includes an AWS CodeBuild build stage that requires access to a database to run integration tests. The developer is using a buildspec.yml file to configure the database connection. Company policy requires automatic rotation of all database credentials.

Which solution will handle the database credentials MOST securely?

- A. Retrieve the credentials from variables that are hardcoded in the buildspec.yml file. Configure an AWS Lambda function to rotate the credentials.
- B. Retrieve the credentials from an environment variable that is linked to a SecureString parameter in AWS Systems Manager Parameter Store. Configure Parameter Store for automatic rotation.
- C. Retrieve the credentials from an environment variable that is linked to an AWS Secrets Manager secret. Configure Secrets Manager for automatic rotation.
- D. Retrieve the credentials from an environment variable that contains the connection string in plaintext. Configure an Amazon EventBridge event to rotate the credentials.

Correct Answer: C

Community vote distribution

C (100%)

65703c1 2 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 2 times

rimaSamir 6 months ago

Answer is C as CodeBuild already supports Secret Manager

upvoted 2 times

Gold07 10 months, 1 week ago

c is the correct answer

upvoted 4 times

cmonthatsme 12 months ago

Selected Answer: C

Secure + Rotation are key words for Secrets Manager

upvoted 4 times

Parsons 12 months ago

Selected Answer: C

C is correct.

Explanation: "requires automatic rotation of all database credentials" => "Secrets Manager for automatic rotation."

With the Systems Manager Parameter Store, you have to do that manually.

upvoted 4 times

cloudenthusiast 12 months ago

Selected Answer: C

Because configure Secrets Manager for automatic rotation

upvoted 2 times

A company is developing a serverless multi-tier application on AWS. The company will build the serverless logic tier by using Amazon API Gateway and AWS Lambda.

While the company builds the logic tier, a developer who works on the frontend of the application must develop integration tests. The tests must cover both positive and negative scenarios, depending on success and error HTTP status codes.

Which solution will meet these requirements with the LEAST effort?

- A. Set up a mock integration for API methods in API Gateway. In the integration request from Method Execution, add simple logic to return either a success or error based on HTTP status code. In the integration response, add messages that correspond to the HTTP status codes.
- B. Create two mock integration resources for API methods in API Gateway. In the integration request, return a success HTTP status code for one resource and an error HTTP status code for the other resource. In the integration response, add messages that correspond to the HTTP status codes.
- C. Create Lambda functions to perform tests. Add simple logic to return either success or error, based on the HTTP status codes. Build an API Gateway Lambda integration. Select appropriate Lambda functions that correspond to the HTTP status codes.
- D. Create a Lambda function to perform tests. Add simple logic to return either success or error-based HTTP status codes. Create a mock integration in API Gateway. Select the Lambda function that corresponds to the HTTP status codes.

Correct Answer: A

Community vote distribution

A (84%)

B (16%)

🗨️ 👤 **Parsons** Highly Voted 👍 12 months ago

Selected Answer: A

A is correct (with the LEAST effort)

"API Gateway supports mock integrations for API methods"

"As an API developer, you decide how API Gateway responds to a mock integration request. For this, you configure the method's integration request and integration response to associate a response with a given status code. "

<https://docs.aws.amazon.com/apigateway/latest/developerguide/how-to-mock-integration.html>

upvoted 10 times

🗨️ 👤 **mooncake1** Most Recent 🕒 6 days, 9 hours ago

Selected Answer: A

B - resource should not be used as such..

One Request should fulfill both success and fail scenario. Another resource means it's not the same request

upvoted 1 times

🗨️ 👤 **65703c1** 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ 👤 **SerialIDr** 6 months, 2 weeks ago

Selected Answer: A

This is an efficient solution. Mock integrations in API Gateway allow you to simulate backend logic directly within API Gateway, without the need for an actual backend like Lambda. You can define the behavior and response (including HTTP status codes and messages) directly in API Gateway, making it ideal for quickly developing and testing various scenarios.



upvoted 2 times

🗨️ 👤 **[Removed]** 11 months, 3 weeks ago

Selected Answer: B

The tests must cover both positive and negative scenarios, depending on success and error HTTP status codes.

upvoted 3 times

  **cloudenthusiast** 12 months ago

Selected Answer: A

A because set up a mock integration for API methods in API Gateway with the least effort.

upvoted 3 times

Users are reporting errors in an application. The application consists of several microservices that are deployed on Amazon Elastic Container Service (Amazon ECS) with AWS Fargate.

Which combination of steps should a developer take to fix the errors? (Choose two.)

- A. Deploy AWS X-Ray as a sidecar container to the microservices. Update the task role policy to allow access to the X-Ray API.
- B. Deploy AWS X-Ray as a daemonset to the Fargate cluster. Update the service role policy to allow access to the X-Ray API.
- C. Instrument the application by using the AWS X-Ray SDK. Update the application to use the PutXrayTrace API call to communicate with the X-Ray API.
- D. Instrument the application by using the AWS X-Ray SDK. Update the application to communicate with the X-Ray daemon.
- E. Instrument the ECS task to send the stdout and stderr output to Amazon CloudWatch Logs. Update the task role policy to allow the cloudwatch:PullLogs action.

Correct Answer: A

Community vote distribution



MG1407 Highly Voted 1 year, 5 months ago

AD

A. You can only use X-ray with Fargate as a side car because there is not EC2 image.

D. <https://github.com/aws-samples/aws-xray-fargate>

upvoted 14 times

lamtany 1 year, 4 months ago

With AWS Fargate, there are no EC2 instances to install the X-Ray daemon onto.

However, the X-Ray daemon is actually provided automatically with Fargate - it runs as an additional container alongside the application containers in the task. So there is no need to deploy it as a sidecar.

When using X-Ray with Fargate, you just need to:

Instrument the application code with the X-Ray SDK

The SDK will communicate with the daemon container provided by Fargate

So you're right that there are no EC2 hosts to install daemons on directly. But Fargate handles running the X-Ray daemon automatically as part of the task, eliminating the need for a sidecar. The SDK can communicate with the daemon container transparently.

upvoted 2 times

Nagoracle 1 year, 3 months ago

I agree - AD

<https://github.com/aws-samples/aws-xray-fargate>

upvoted 1 times

vkovilam Most Recent 7 months, 3 weeks ago

A. Deploy AWS X-Ray as a sidecar container to the microservices. Update the task role policy to allow access to the X-Ray API.

D. Instrument the application by using the AWS X-Ray SDK. Update the application to communicate with the X-Ray daemon.

upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: A

AD is the correct answer.

upvoted 1 times

Single 10 months, 1 week ago

Selected Answer: AE

ChatGPT/AE

upvoted 1 times

🗨️ **41eb566** 10 months, 2 weeks ago

Selected Answer: C

CE

C. Instrument the application by using the AWS X-Ray SDK. Update the application to use the PutXrayTrace API call to communicate with the X-Ray API.

This step involves instrumenting the application code using the AWS X-Ray SDK to generate trace data and communicate it to the X-Ray service for analysis.

E. Instrument the ECS task to send the stdout and stderr output to Amazon CloudWatch Logs. Update the task role policy to allow the cloudwatch:PullLogs action.

upvoted 1 times

🗨️ **rrshah83** 1 year, 1 month ago

Selected Answer: A

AC

Fargate cannot have daemon. This rules out B and C. D is distractor.

upvoted 1 times

🗨️ **tqiu654** 1 year, 1 month ago

Selected Answer: A

CHatGpt: AD

upvoted 1 times

🗨️ **ShawnWon** 1 year, 2 months ago

DE

Option A is incorrect because deploying AWS X-Ray as a sidecar container to the microservices is not the common practice for Fargate deployments. Fargate tasks usually run as a single container, and the application is instrumented to communicate with the X-Ray daemon.

Option B is not applicable because deploying AWS X-Ray as a daemonset is a concept related to Kubernetes, not AWS Fargate.

Option C is incorrect because using the AWS X-Ray SDK involves instrumenting the application, but the suggested approach is to communicate with the X-Ray daemon rather than directly calling the X-Ray API.

upvoted 3 times

🗨️ **Passexam4sure_com** 1 year, 3 months ago

Selected Answer: D

Instrument the application by using the AWS X-Ray SDK. Update the application to communicate with the X-Ray daemon

upvoted 1 times

🗨️ **Claire_KMT** 1 year, 3 months ago

D. Instrument the application by using the AWS X-Ray SDK. Update the application to communicate with the X-Ray daemon.

E. Instrument the ECS task to send the stdout and stderr output to Amazon CloudWatch Logs. Update the task role policy to allow the cloudwatch:PullLogs action.

upvoted 1 times

🗨️ **fossil123** 1 year, 4 months ago

Selected Answer: A

AD is correct.

A - X-Ray container as a "Side car" in ECS/Fargate cluster

D - Instrument the application using the AWS X-Ray SDK to collect telemetry data.

upvoted 3 times

🗨️ **love777** 1 year, 5 months ago

Selected Answer: D

D and E

Option D:

Instrumenting the application using the AWS X-Ray SDK is essential for collecting traces and telemetry data. The X-Ray SDK helps you identify

bottlenecks, errors, and other issues within your microservices.

Communicating with the X-Ray daemon allows your microservices to send trace data to X-Ray for analysis and visualization. This requires minimal configuration and is efficient for capturing and analyzing traces.

Option E:

Instrumenting the ECS task to send the application's standard output (stdout) and standard error (stderr) logs to Amazon CloudWatch Logs provides visibility into the application's behavior, errors, and issues.

Updating the task role policy to allow the cloudwatch:PullLogs action ensures that the ECS task has the necessary permissions to access and send logs to CloudWatch Logs.

upvoted 3 times

 **AWSdeveloper08** 1 year, 5 months ago

Selected Answer: C

Answer is CE

To diagnose and fix errors in an application deployed on Amazon ECS with AWS Fargate using AWS X-Ray, you should take the following steps:

C. Instrument the application by using the AWS X-Ray SDK. Update the application to use the PutXrayTrace API call to communicate with the X-Ray API.

Instrumenting the application using the AWS X-Ray SDK allows you to capture traces and data about requests as they flow through your application's components.

E. Instrument the ECS task to send the stdout and stderr output to Amazon CloudWatch Logs. Update the task role policy to allow the cloudwatch:PullLogs action.

This step will help you capture logs from your microservices, which can provide additional insights into the errors and issues occurring within the application.

upvoted 1 times

A developer is creating an application for a company. The application needs to read the file doc.txt that is placed in the root folder of an Amazon S3 bucket that is named DOC-EXAMPLE-BUCKET. The company's security team requires the principle of least privilege to be applied to the application's IAM policy.

Which IAM policy statement will meet these security requirements?

- A.

```
{
  "Action": [
    "s3:GetObject"
  ],
  "Effect": "Allow",
  "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET/doc.txt"
}
```
- B.

```
{
  "Action": [
    "s3:*"
  ],
  "Effect": "Allow",
  "Resource": "*"
}
```
- C.

```
{
  "Action": [
    "s3:GetObject"
  ],
  "Effect": "Allow",
  "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
}
```
- D.

```
{
  "Action": [
    "s3:*"
  ],
  "Effect": "Allow",
  "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET/doc.txt"
}
```

Correct Answer: A

Community vote distribution

A (100%)

 **Gadu** Highly Voted 11 months, 4 weeks ago

Selected Answer: A

Only read permission for the file
upvoted 7 times

 **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: A

A is the correct answer.
upvoted 1 times

 **cmonthatsme** 12 months ago

Selected Answer: A

Only allow to get this one file. A
upvoted 4 times


A company has an application that uses AWS CodePipeline to automate its continuous integration and continuous delivery (CI/CD) workflow. The application uses AWS CodeCommit for version control. A developer who was working on one of the tasks did not pull the most recent changes from the main branch. A week later, the developer noticed merge conflicts.

How can the developer resolve the merge conflicts in the developer's branch with the LEAST development effort?

- A. Clone the repository. Create a new branch. Update the branch with the changes.
- B. Create a new branch. Apply the changes from the previous branch.
- C. Use the Commit Visualizer view to compare the commits when a feature was added. Fix the merge conflicts.
- D. Stop the pull from the main branch to the feature branch. Rebase the feature branch from the main branch.

Correct Answer: D

Community vote distribution



love777 Highly Voted 1 year, 5 months ago

Selected Answer: D

Option D is the best approach for resolving the merge conflicts with minimal development effort. Here's how it works:

Stop Pull from Main: By stopping the pull from the main branch to the feature branch, the developer can prevent the introduction of new conflicts while they are resolving the existing ones.

Rebase the Feature Branch: After stopping the pull, the developer can rebase the feature branch onto the main branch. This essentially replays the feature branch's changes on top of the main branch's latest changes. This allows the developer to resolve conflicts one commit at a time, addressing any conflicts that arise from the difference between the feature branch and the main branch.

upvoted 10 times

9d8dd9c Most Recent 3 months ago

Most people say D, but option D does not cover the fixing of conflict, but option C at least says "fix the conflicts" so this is a complete answer, no?

upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

maurice2005 11 months ago

Selected Answer: C

How on earth non-visualize way is easier? And I've never seen a rebase that happened commit-by-commit! Resolving the merge should have happened in one go, merge or rebase. just rebase appear(!) more isolated. (which in practice it's not!). Also rebase is cleaner but the effort is even more since there is a level of isolation!

The only point here is visualizing which makes it easier.

upvoted 2 times

SerialiDr 1 year ago

Selected Answer: D

D. Stop the pull from the main branch to the feature branch. Rebase the feature branch from the main branch: Rebasing the feature branch from the main branch is an effective way to resolve merge conflicts. This approach involves updating the feature branch with the latest changes from the main branch and then applying the feature branch's changes on top of it. Rebasing can simplify the process of resolving conflicts and is generally less effort-intensive compared to creating new branches and transferring changes.

C. Use the Commit Visualizer view to compare the commits when a feature was added. Fix the merge conflicts: Using tools like Commit Visualizer to understand the changes and conflicts can be helpful. However, this step alone doesn't resolve the conflicts. The developer still needs to manually resolve the conflicts in the code.

upvoted 1 times

🗨️ 👤 **Passexam4sure_com** 1 year, 3 months ago

D

D. Stop the pull from the main branch to the feature branch. Rebase the feature branch from the main branch.

upvoted 1 times

🗨️ 👤 **Claire_KMT** 1 year, 3 months ago

D. Stop the pull from the main branch to the feature branch. Rebase the feature branch from the main branch.

upvoted 1 times

🗨️ 👤 **lamtany** 1 year, 4 months ago

Selected Answer: D

Rebasing the feature branch from the main branch would apply the changes from the main branch directly onto the feature branch, effectively bringing it up to date. This would resolve the conflicts in a way that minimizes manual effort.

upvoted 3 times

🗨️ 👤 **[Removed]** 1 year, 5 months ago

Selected Answer: D

Option D is the best approach for resolving the merge conflicts

upvoted 2 times

🗨️ 👤 **[Removed]** 1 year, 5 months ago

Selected Answer: D

Using the git rebase command to rebase a repository changes the history of a repository, which might cause commits to appear out of order.

<https://docs.aws.amazon.com/codecommit/latest/userguide/how-to-view-commit-details.html>

upvoted 1 times

🗨️ 👤 **AWSdeveloper08** 1 year, 5 months ago

Selected Answer: C

Comparing commits in the Commit Visualizer view can provide a clear overview of the changes made over time and aid in understanding the context of the conflicts. This approach can help you pinpoint where conflicts arose and assist you in making informed decisions about how to resolve them.

upvoted 4 times

🗨️ 👤 **worseforwear** 1 year, 5 months ago

Selected Answer: C

Answer D won't fix the problem

upvoted 1 times

🗨️ 👤 **Cerakoted** 1 year, 3 months ago

I think C would take huge development effort

upvoted 1 times

🗨️ 👤 **maurice2005** 11 months ago

because visualizing make it harder? You have to fix the conflict anyway! rebase or merge. in both resolve the conflict will happened in one go (unlike the comments I see which they say rebase is commit by commit). I don't think those who pick rebase ever used it before in practice!

upvoted 1 times

A developer wants to add request validation to a production environment Amazon API Gateway API. The developer needs to test the changes before the API is deployed to the production environment. For the test, the developer will send test requests to the API through a testing tool.

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

- A. Export the existing API to an OpenAPI file. Create a new API. Import the OpenAPI file. Modify the new API to add request validation. Perform the tests. Modify the existing API to add request validation. Deploy the existing API to production.
- B. Modify the existing API to add request validation. Deploy the updated API to a new API Gateway stage. Perform the tests. Deploy the updated API to the API Gateway production stage.
- C. Create a new API. Add the necessary resources and methods, including new request validation. Perform the tests. Modify the existing API to add request validation. Deploy the existing API to production
- D. Clone the existing API. Modify the new API to add request validation. Perform the tests. Modify the existing API to add request validation. Deploy the existing API to production.

Correct Answer: B

Community vote distribution

B (100%)

 **AWSdeveloper08** Highly Voted 11 months, 3 weeks ago

Selected Answer: B

In this option, you are making changes directly to the existing API, adding request validation. Then, you deploy the updated API to a new API Gateway stage, which allows you to test the changes without affecting the production environment. After performing the tests and ensuring everything works as expected, you can then deploy the updated API to the production stage, thus minimizing operational overhead.


upvoted 12 times

 **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

 **SerialiDr** 6 months, 2 weeks ago

Selected Answer: B

B. Modify the existing API to add request validation. Deploy the updated API to a new API Gateway stage. Perform the tests. Deploy the updated API to the API Gateway production stage: This is a more streamlined approach. By deploying the updated API to a new stage, the developer can test the changes in an environment that closely mirrors production without affecting the current production traffic. Once testing is complete, the changes can be deployed to the production stage. This approach minimizes operational overhead.

upvoted 2 times

 **imyashkale** 10 months, 3 weeks ago

Selected Answer: B

It looks Correct

upvoted 2 times

An online food company provides an Amazon API Gateway HTTP API to receive orders for partners. The API is integrated with an AWS Lambda function. The Lambda function stores the orders in an Amazon DynamoDB table.

The company expects to onboard additional partners. Some of the partners require additional Lambda functions to receive orders. The company has created an Amazon S3 bucket. The company needs to store all orders and updates in the S3 bucket for future analysis.




How can the developer ensure that all orders and updates are stored to Amazon S3 with the LEAST development effort?

- A. Create a new Lambda function and a new API Gateway API endpoint. Configure the new Lambda function to write to the S3 bucket. Modify the original Lambda function to post updates to the new API endpoint.
- B. Use Amazon Kinesis Data Streams to create a new data stream. Modify the Lambda function to publish orders to the data stream. Configure the data stream to write to the S3 bucket.
- C. Enable DynamoDB Streams on the DynamoDB table. Create a new Lambda function. Associate the stream's Amazon Resource Name (ARN) with the Lambda function. Configure the Lambda function to write to the S3 bucket as records appear in the table's stream.
- D. Modify the Lambda function to publish to a new Amazon Simple Notification Service (Amazon SNS) topic as the Lambda function receives orders. Subscribe a new Lambda function to the topic. Configure the new Lambda function to write to the S3 bucket as updates come through the topic.

Correct Answer: C

Community vote distribution

C (100%)

  **AWSdeveloper08** Highly Voted  11 months, 3 weeks ago

Selected Answer: C

By enabling DynamoDB Streams on the DynamoDB table, you can capture changes (orders and updates) to the table. Whenever a new order or an update is made to the table, a stream record is generated. You can then create a new Lambda function, associate the stream's ARN with this Lambda function, and configure it to write the stream records (orders and updates) to the S3 bucket. This approach leverages built-in features of DynamoDB and Lambda, minimizing the development effort required to achieve the desired outcome.



upvoted 9 times

  **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: C

C is the correct answer.


upvoted 1 times

  **SerialIDr** 6 months, 2 weeks ago

Selected Answer: C

This is a streamlined and effective approach. Enabling DynamoDB Streams captures modifications to the DynamoDB table (such as new orders) and triggers a new Lambda function. This function can then write these changes to the S3 bucket. This approach requires minimal changes to the existing setup and leverages the integration between DynamoDB Streams and Lambda.

upvoted 2 times

  **Dushank** 10 months, 3 weeks ago

Selected Answer: C

Enabling DynamoDB Streams on the existing DynamoDB table and associating a new Lambda function to it would be a straightforward way to capture all changes (new orders and updates) in the DynamoDB table. The new Lambda function would automatically be triggered when a new record appears in the table's stream and could be configured to write this data to the S3 bucket. This is likely the least effort-intensive approach for meeting the requirement.

upvoted 3 times

A company's website runs on an Amazon EC2 instance and uses Auto Scaling to scale the environment during peak times. Website users across the world are experiencing high latency due to static content on the EC2 instance, even during non-peak hours.


Which combination of steps will resolve the latency issue? (Choose two.)

- A. Double the Auto Scaling group's maximum number of servers.
- B. Host the application code on AWS Lambda.
- C. Scale vertically by resizing the EC2 instances.
- D. Create an Amazon CloudFront distribution to cache the static content.
- E. Store the application's static content in Amazon S3.

Correct Answer: DE

Community vote distribution


DE (100%)

 **65703c1** 2 months, 1 week ago

Selected Answer: DE

DE is the correct answer.

upvoted 2 times

 **SerialiDr** 6 months, 2 weeks ago

Selected Answer: DE

D. Create an Amazon CloudFront distribution to cache the static content: This is an effective solution. Amazon CloudFront is a content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency and high transfer speeds. Using CloudFront to cache static content closer to users can significantly reduce latency.

E. Store the application's static content in Amazon S3: This is another effective solution. Amazon S3 can serve as a highly durable and scalable storage solution for static content. When combined with Amazon CloudFront, it provides an efficient way to manage and deliver static content with reduced latency.

The combination of steps that will best resolve the latency issue is:

D. Create an Amazon CloudFront distribution to cache the static content.

E. Store the application's static content in Amazon S3.

upvoted 3 times

 **Digo30sp** 10 months ago

Selected Answer: DE

Option (D), creating an Amazon CloudFront distribution to cache static content, is the most recommended solution. CloudFront is a global content delivery network (CDN) that can cache static content on servers distributed around the world. This can help significantly reduce latency for users around the world. Option (E), storing your application's static content in Amazon S3, can also help reduce latency. S3 is a high-performance object storage service that can be used to store static content.

upvoted 4 times


A company has an Amazon S3 bucket containing premier content that it intends to make available to only paid subscribers of its website. The S3 bucket currently has default permissions of all objects being private to prevent inadvertent exposure of the premier content to non-paying website visitors.

How can the company limit the ability to download a premier content file in the S3 bucket to paid subscribers only?

- A. Apply a bucket policy that allows anonymous users to download the content from the S3 bucket.
- B. Generate a pre-signed object URL for the premier content file when a paid subscriber requests a download.
- C. Add a bucket policy that requires multi-factor authentication for requests to access the S3 bucket objects.
- D. Enable server-side encryption on the S3 bucket for data protection against the non-paying website visitors.

Correct Answer: B

Community vote distribution



🗨️ **BrainFried** 6 months, 1 week ago

Selected Answer: C

The answer is NOT B.

The question states "limit the ABILITY TO DOWNLOAD a file in the S3 bucket to PAID SUBS ONLY"

If you choose B, it means that a paid sub can request the URL and then SHARE IT with non-paying customers. This will mean non-paying customers can DOWNLOAD the file using the URL.

The answer should be C - this enforces the user who downloads the file is a paying customer.

upvoted 1 times

🗨️ **65703c1** 8 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 2 times

🗨️ **SerialiDr** 1 year ago

Selected Answer: B

B. Generate a pre-signed object URL for the premier content file when a paid subscriber requests a download: This is the most appropriate solution. A pre-signed URL grants temporary access to a private object stored in S3. The URL can be generated programmatically, and its validity can be limited to a short duration. This approach allows only those who have been provided with the URL (paid subscribers, in this case) to download the specific content.

upvoted 3 times

🗨️ **Digo30sp** 1 year, 3 months ago

Selected Answer: B

The correct answer is (B).

By generating a pre-signed object URL for the main content file when a paid subscriber requests a download, the company can control who can download the file. The pre-signed object URL will be valid for a limited period of time and can only be used by the paid subscriber who requested the download.

upvoted 4 times

A developer is creating an AWS Lambda function that searches for items from an Amazon DynamoDB table that contains customer contact information. The DynamoDB table items have the customer's email_address as the partition key and additional properties such as customer_type, name and job_title.

The Lambda function runs whenever a user types a new character into the customer_type text input. The developer wants the search to return partial matches of all the email_address property of a particular customer_type. The developer does not want to recreate the DynamoDB table.

What should the developer do to meet these requirements?

- A. Add a global secondary index (GSI) to the DynamoDB table with customer_type as the partition key and email_address as the sort key. Perform a query operation on the GSI by using the begins_with key condition expression with the email_address property.
- B. Add a global secondary index (GSI) to the DynamoDB table with email_address as the partition key and customer_type as the sort key. Perform a query operation on the GSI by using the begins_with key condition expression with the email_address property.
- C. Add a local secondary index (LSI) to the DynamoDB table with customer_type as the partition key and email_address as the sort key. Perform a query operation on the LSI by using the begins_with key condition expression with the email_address property.
- D. Add a local secondary index (LSI) to the DynamoDB table with job_title as the partition key and email_address as the sort key. Perform a query operation on the LSI by using the begins_with key condition expression with the email_address property.

Correct Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **Saurabh04** 5 months, 3 weeks ago

Option B (Global Secondary Index with email_address as Partition Key):

Add a global secondary index (GSI) to the DynamoDB table.

Set email_address as the partition key and customer_type as the sort key for the GSI.

Perform a query operation on the GSI using the begins_with key condition expression with the email_address property.

upvoted 1 times

🗳️ 👤 **65703c1** 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 3 times

🗳️ 👤 **Examenee** 11 months, 4 weeks ago

Selected Answer: A

Only global secondary indices can be added after a table has been created.

upvoted 3 times

🗳️ 👤 **SerialiDr** 1 year ago

Selected Answer: A

A. Add a global secondary index (GSI) to the DynamoDB table with customer_type as the partition key and email_address as the sort key. Perform a query operation on the GSI by using the begins_with key condition expression with the email_address property: This approach is correct. By creating a GSI with customer_type as the partition key and email_address as the sort key, the developer can efficiently query items based on customer_type. The begins_with condition can be applied to the sort key (email_address) in the GSI, allowing for searches that return partial matches.

upvoted 1 times

🗳️ 👤 **RamyaMunipala** 1 year ago

Selected Answer: A

A is correct

upvoted 1 times

🗳️ 👤 **Jing2023** 1 year, 3 months ago

A is correct

upvoted 1 times

  **Patel_ajay745** 1 year, 3 months ago

A

Add a global secondary index (GSI) to the DynamoDB table with customer_type as the partition key and email_address as the sort key. Perform a query operation on the GSI by using the begins_with key condition expression with the email_address property.

upvoted 2 times

  **Digo30sp** 1 year, 3 months ago

Selected Answer: A

The correct answer is (A).

By adding a global secondary index (GSI) to the DynamoDB table with customer_type as the partition key and email_address as the sort key, the developer can perform a query operation on the GSI using the Begins_with key condition expression with the email_address property. This will return partial matches of all email_address properties of a specific customer_type.

upvoted 4 times

A developer is building an application that uses AWS API Gateway APIs, AWS Lambda functions, and AWS DynamoDB tables. The developer uses the AWS Serverless Application Model (AWS SAM) to build and run serverless applications on AWS. Each time the developer pushes changes for only to the Lambda functions, all the artifacts in the application are rebuilt.

The developer wants to implement AWS SAM Accelerate by running a command to only redeploy the Lambda functions that have changed.

Which command will meet these requirements?

- A. `sam deploy --force-upload`
- B. `sam deploy --no-execute-changeset`
- C. `sam package`
- D. `sam sync --watch`

Correct Answer: D

Community vote distribution

D (100%)

 **Digo30sp** Highly Voted 10 months ago

Selected Answer: D

The correct answer is (D).

The `sam sync --watch` command will only deploy the Lambda functions that have changed. This command uses AWS SAM Accelerate to compare the local versions of your Lambda functions to the versions deployed in AWS. If there are differences, the command deploys only the changed Lambda functions.


upvoted 8 times

 **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: D

D is the correct answer.


upvoted 1 times

 **hayjaykay** 5 months, 2 weeks ago

Correct answer is B.

To deploy only the Lambda functions that have changed using AWS SAM Accelerate, the developer can use the `sam deploy --no-execute-changeset` command. This command will create an AWS CloudFormation change set without executing it, allowing the developer to preview the changes before deploying.

upvoted 1 times


 **SerialiDr** 6 months, 2 weeks ago

Selected Answer: D

D. `sam sync --watch`: This command is a part of SAM Accelerate and is used for rapid iterative development. When run, it watches for changes in the source files of your Lambda functions and APIs and deploys only those changes, rather than redeploying the entire stack. This greatly speeds up the deployment process during development.

Therefore, to implement AWS SAM Accelerate and only redeploy the Lambda functions that have changed, the developer should use `sam sync --watch`. This command aligns with the goal of deploying changes rapidly and efficiently, focusing only on the components that have been modified.

upvoted 1 times

 **dilleman** 9 months, 3 weeks ago

Selected Answer: D

D is correct

upvoted 1 times

A developer is building an application that gives users the ability to view bank accounts from multiple sources in a single dashboard. The developer has automated the process to retrieve API credentials for these sources. The process invokes an AWS Lambda function that is associated with an AWS CloudFormation custom resource.

The developer wants a solution that will store the API credentials with minimal operational overhead.

Which solution will meet these requirements in the MOST secure way?

- A. Add an AWS Secrets Manager GenerateSecretString resource to the CloudFormation template. Set the value to reference new credentials for the CloudFormation resource.
- B. Use the AWS SDK ssm:PutParameter operation in the Lambda function from the existing custom resource to store the credentials as a parameter. Set the parameter value to reference the new credentials. Set the parameter type to SecureString.
- C. Add an AWS Systems Manager Parameter Store resource to the CloudFormation template. Set the CloudFormation resource value to reference the new credentials. Set the resource NoEcho attribute to true.
- D. Use the AWS SDK ssm:PutParameter operation in the Lambda function from the existing custom resource to store the credentials as a parameter. Set the parameter value to reference the new credentials. Set the parameter NoEcho attribute to true.

Correct Answer: B

Community vote distribution



🗨️ 👤 **Jing2023** Highly Voted 1 year, 3 months ago

Answer is B

A is not correct as the requirement asked to store API credentials, GenerateSecretString will create a random string as password.

C the API credential will be retrieved by the Lambda function, it is un-available to the template.

D no echo is a attribute of cloud formation template.

upvoted 17 times

🗨️ 👤 **Digo30sp** Highly Voted 1 year, 3 months ago

Selected Answer: D

The correct answer is (D).

Solution (D) is the most secure because it stores the API credentials in AWS Secrets Manager, which is a managed service that provides secure, policy-controlled storage for secrets. The parameter's NoEcho attribute prevents the parameter value from being displayed in the console or request history.

upvoted 7 times

🗨️ 👤 **ShakthiGCP** Most Recent 2 months, 4 weeks ago

Selected Answer: B

SecureString parameters are encrypted both when stored in the Parameter Store (at rest) and while being transmitted (in transit) using AWS KMS (Key Management Service). This means that even if someone were to gain unauthorized access to the Parameter Store's underlying storage, they wouldn't be able to easily read the parameter's value.

upvoted 1 times

🗨️ 👤 **Anandesh** 6 months, 2 weeks ago

Selected Answer: D

https://docs.aws.amazon.com/systems-manager/latest/APIReference/API_PutParameter.html

upvoted 1 times

🗨️ 👤 **65703c1** 8 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times


🗨️ 👤 **Melisa202401** 9 months, 4 weeks ago

Selected Answer: C

I choose C, not choose A due to minimal cost


I dont understand why most of you choose B

upvoted 2 times

 **ShakthiGCP** 2 months, 4 weeks ago

NoEcho is not very secure compared to SecureString which encrypts the sensitive data at rest and transit.

upvoted 1 times

 **KarBiswa** 10 months, 2 weeks ago

its B only

<https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-parameter-store.html#:~:text=You%20can%20also%20use%20SecureString%20parameters%20with%20other%20AWS%20services.%20In%20the%20following%20exam>

upvoted 2 times

 **KarBiswa** 11 months, 1 week ago

Selected Answer: A

I will got with A.

Because <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/parameters-section-structure.html> nullifying the B&D. Justifying A

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/dynamic-references.html>

upvoted 1 times

 **KarBiswa** 10 months, 2 weeks ago

Reverting the Option to B

upvoted 1 times

 **KillThemWithKindness** 11 months, 2 weeks ago

Selected Answer: B

The solution that will meet the requirements is to use the AWS SDK ssm PutParameter operation in the Lambda function from the existing custom resource to store the credentials as a parameter. Set the parameter value to reference the new credentials. Set the parameter type to SecureString. This way, the developer can store the API credentials with minimal operational overhead, as AWS Systems Manager Parameter Store provides secure and scalable storage for configuration data. The SecureString parameter type encrypts the parameter value with AWS Key Management Service (AWS KMS). The other options either involve adding additional resources to the CloudFormation template, which increases complexity and cost, or do not encrypt the parameter value, which reduces security.

upvoted 2 times

 **SerialiDr** 1 year ago

Selected Answer: B

B. Use the AWS SDK ssm:PutParameter operation in the Lambda function from the existing custom resource to store the credentials as a parameter. Set the parameter value to reference the new credentials. Set the parameter type to SecureString: This is a secure and operationally efficient solution. AWS Systems Manager Parameter Store can securely store parameters as SecureString, which encrypts the parameter value. The ssm:PutParameter operation can be used within the Lambda function to store the credentials directly after retrieval, minimizing operational overhead.

upvoted 1 times

 **Snape** 1 year ago

Selected Answer: B

Answer is B

upvoted 1 times

 **rrshah83** 1 year, 1 month ago

Selected Answer: B

noecho is CF feature, not ssm param store

upvoted 1 times

 **Certified101** 1 year, 1 month ago

Selected Answer: B

Agree with B - D will be stored in plain text, this is credentials so should be secure string

upvoted 2 times

 **kaes** 1 year, 2 months ago

Selected Answer: D

ANS: D

NoEcho <https://github.com/aws-cloudformation/cloudformation-coverage-roadmap/issues/82#issuecomment-517704282>

upvoted 3 times

🗨️ 👤 **kaes** 1 year, 2 months ago

ANS: D

NoEcho <https://github.com/aws-cloudformation/cloudformation-coverage-roadmap/issues/82#issuecomment-517704282>

upvoted 1 times

🗨️ 👤 **ut18** 1 year, 3 months ago

Is B the correct answer?

SecureString isn't currently supported for AWS CloudFormation templates.

https://docs.aws.amazon.com/systems-manager/latest/APIReference/API_PutParameter.html

upvoted 2 times

🗨️ 👤 **Bolu_Jay** 1 year, 3 months ago

Answer is A

AWS Secrets Manager is specifically designed for securely storing sensitive information like API credentials, database passwords, and other secrets

upvoted 5 times

A developer is trying to get data from an Amazon DynamoDB table called demoman-table. The developer configured the AWS CLI to use a specific IAM user's credentials and ran the following command:

```
aws dynamodb get-item --table-name demoman-table --key '{"id": {"N": "1993"}}
```

The command returned errors and no rows were returned.


What is the MOST likely cause of these issues?

- A. The command is incorrect; it should be rewritten to use put-item with a string argument.
- B. The developer needs to log a ticket with AWS Support to enable access to the demoman-table.
- C. Amazon DynamoDB cannot be accessed from the AWS CLI and needs to be called via the REST API.
- D. The IAM user needs an associated policy with read access to demoman-table.

Correct Answer: D

Community vote distribution


D (100%)

 **albert_kuo** 4 months, 1 week ago

Selected Answer: D

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "dynamodb:GetItem",
      "Resource": "arn:aws:dynamodb:us-east-1:123456789012:table/demoman-table"
    }
  ]
}
```

upvoted 1 times

 **65703c1** 8 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

 **Jing2023** 1 year, 3 months ago

Selected Answer: D

D is correct

upvoted 3 times

 **dilleman** 1 year, 3 months ago

Selected Answer: D

D is correct

upvoted 1 times

 **Digo30sp** 1 year, 3 months ago

Selected Answer: D

The correct answer is (D).

The command is correct and the demoman table exists. The most likely issue is that the IAM user does not have a policy associated with read access to the demoman table.

To resolve the issue, the developer must add a policy to the IAM user that grants read access to the demoman table.

upvoted 4 times

An organization is using Amazon CloudFront to ensure that its users experience low-latency access to its web application. The organization has identified a need to encrypt all traffic between users and CloudFront, and all traffic between CloudFront and the web application.

How can these requirements be met? (Choose two.)

- A. Use AWS KMS to encrypt traffic between CloudFront and the web application.
- B. Set the Origin Protocol Policy to "HTTPS Only".
- C. Set the Origin's HTTP Port to 443.
- D. Set the Viewer Protocol Policy to "HTTPS Only" or "Redirect HTTP to HTTPS".
- E. Enable the CloudFront option Restrict Viewer Access.

Correct Answer: BD

Community vote distribution

BD (100%)

 **SerialiDr** Highly Voted 6 months, 2 weeks ago

Selected Answer: BD

B. Set the Origin Protocol Policy to "HTTPS Only": This setting ensures that all traffic between CloudFront and the web application (origin) is encrypted. By setting the Origin Protocol Policy to "HTTPS Only," CloudFront will only connect to the origin over HTTPS, ensuring encryption of data in transit.

D. Set the Viewer Protocol Policy to "HTTPS Only" or "Redirect HTTP to HTTPS": This setting is crucial for ensuring that all traffic between the users (viewers) and CloudFront is encrypted. By setting the Viewer Protocol Policy to "HTTPS Only" or "Redirect HTTP to HTTPS," CloudFront ensures that user requests are either only served over HTTPS or automatically redirected from HTTP to HTTPS.

upvoted 5 times

 **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: BD

BD is the correct answer.

upvoted 1 times


 **Jeff1719** 8 months ago

Selected Answer: BD

BD: Protocol and Viewer protocol policy, see

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/distribution-web-values-specify.html>

upvoted 2 times

 **dilleman** 9 months, 3 weeks ago

Selected Answer: BD

B and D are the correct ones.

B: Setting the Origin Protocol Policy to "HTTPS Only" ensures that CloudFront always uses HTTPS to connect to the origin, which is the web application in this scenario.

D: Setting the Viewer Protocol Policy to "HTTPS Only" ensures that CloudFront will only serve requests over HTTPS. Setting it to "Redirect HTTP to HTTPS" ensures that any HTTP request from viewers is redirected to HTTPS.

upvoted 4 times

 **Digo30sp** 10 months ago

Selected Answer: BD

The correct answers are (B) and (D).

To meet the requirement to encrypt all traffic between users and CloudFront, your organization must set the Viewer Protocol Policy to "HTTPS Only" or "Redirect HTTP to HTTPS". This will force users to use HTTPS to connect to CloudFront.

To meet the requirement to encrypt all traffic between CloudFront and the web application, your organization must set the Origin Protocol Policy to "HTTPS Only". This will force CloudFront to use HTTPS to connect to the web application.

upvoted 3 times

A developer is planning to migrate on-premises company data to Amazon S3. The data must be encrypted, and the encryption keys must support automatic annual rotation. The company must use AWS Key Management Service (AWS KMS) to encrypt the data.

Which type of keys should the developer use to meet these requirements?

- A. Amazon S3 managed keys
- B. Symmetric customer managed keys with key material that is generated by AWS
- C. Asymmetric customer managed keys with key material that is generated by AWS
- D. Symmetric customer managed keys with imported key material

Correct Answer: B

Community vote distribution

B (71%)

A (29%)

 **PrakashM14** Highly Voted 9 months, 3 weeks ago

Selected Answer: B

Asymmetric keys (option C) are typically used for different use cases, such as digital signatures and key pairs, and may not be as suitable for automatic rotation in the described scenario.

Imported key material (option D) means that you bring your own key material, and AWS KMS doesn't support automatic rotation for such keys.

Amazon S3 managed keys (option A) are used specifically for Amazon S3 and don't support automatic rotation.

so, option B is correct

upvoted 13 times

 **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: B

B is the correct answer.


upvoted 1 times

 **SerialiDr** 5 months ago

Selected Answer: B

This option allows for automatic rotation of the keys, aligning with AWS best practices for key management and security. AWS KMS supports key rotation, which can be configured to occur automatically on an annual basis for customer managed keys. This ensures that data remains encrypted with a key that is periodically rotated, enhancing the security posture of the data stored in Amazon S3.

upvoted 2 times

 **KarBiswa** 5 months, 1 week ago

Selected Answer: B

<https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html> Its a symmetric key rotation

upvoted 1 times


 **koniczny69** 6 months ago

Selected Answer: A

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingServerSideEncryption.html>

Server-side encryption protects data at rest. Amazon S3 encrypts each object with a unique key. As an additional safeguard, it encrypts the key itself with a key that it rotates regularly. Amazon S3 server-side encryption uses 256-bit Advanced Encryption Standard Galois/Counter Mode (AES-GCM) to encrypt all uploaded objects.

upvoted 3 times

 **SerialiDr** 6 months, 2 weeks ago

Selected Answer: B

B. Symmetric customer managed keys with key material that is generated by AWS: This option allows the developer to create and manage their own encryption keys in AWS KMS, with AWS generating the key material. AWS KMS supports automatic rotation of customer managed keys. You

can configure the key to rotate automatically once per year.

upvoted 2 times

🗨️ **Certified101** 7 months, 3 weeks ago

Selected Answer: B

B is correct, it must use KMS

upvoted 1 times

🗨️ **ShawnWon** 8 months, 2 weeks ago

Option A (Amazon S3 managed keys) does not involve using AWS Key Management Service (AWS KMS) directly. Instead, it relies on Amazon S3 to manage the keys for server-side encryption. If the requirement is specifically to use AWS KMS for encryption, then Option A would not meet that requirement.

upvoted 1 times

🗨️ **wonder_man** 9 months, 1 week ago

Selected Answer: B

Only this option supports AWS KMS with the key rotation

upvoted 1 times

🗨️ **PrakashM14** 9 months, 3 weeks ago

Asymmetric keys (option C) are typically used for different use cases, such as digital signatures and key pairs, and may not be as suitable for automatic rotation in the described scenario.

Imported key material (option D) means that you bring your own key material, and AWS KMS doesn't support automatic rotation for such keys.

Amazon S3 managed keys (option A) are used specifically for Amazon S3 and don't support automatic rotation.

so, option B is correct

upvoted 1 times

🗨️ **dilleman** 9 months, 3 weeks ago

Selected Answer: A

A: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingServerSideEncryption.html>

upvoted 2 times

🗨️ **Digo30sp** 10 months ago

Selected Answer: A

A) Amazon S3 Managed Keys

https://docs.aws.amazon.com/pt_br/AmazonS3/latest/userguide/serv-side-encryption.html

upvoted 3 times


A team of developers is using an AWS CodePipeline pipeline as a continuous integration and continuous delivery (CI/CD) mechanism for a web application. A developer has written unit tests to programmatically test the functionality of the application code. The unit tests produce a test report that shows the results of each individual check. The developer now wants to run these tests automatically during the CI/CD process.

Which solution will meet this requirement with the LEAST operational effort?

- A. Write a Git pre-commit hook that runs the tests before every commit. Ensure that each developer who is working on the project has the pre-commit hook installed locally. Review the test report and resolve any issues before pushing changes to AWS CodeCommit.
- B. Add a new stage to the pipeline. Use AWS CodeBuild as the provider. Add the new stage after the stage that deploys code revisions to the test environment. Write a buildspec that fails the CodeBuild stage if any test does not pass. Use the test reports feature of CodeBuild to integrate the report with the CodeBuild console. View the test results in CodeBuild. Resolve any issues.
- C. Add a new stage to the pipeline. Use AWS CodeBuild as the provider. Add the new stage before the stage that deploys code revisions to the test environment. Write a buildspec that fails the CodeBuild stage if any test does not pass. Use the test reports feature of CodeBuild to integrate the report with the CodeBuild console. View the test results in CodeBuild. Resolve any issues.
- D. Add a new stage to the pipeline. Use Jenkins as the provider. Configure CodePipeline to use Jenkins to run the unit tests. Write a Jenkinsfile that fails the stage if any test does not pass. Use the test report plugin for Jenkins to integrate the report with the Jenkins dashboard. View the test results in Jenkins. Resolve any issues.

Correct Answer: C

Community vote distribution



SerialIDr Highly Voted 1 year ago

Selected Answer: C

This is the most efficient and integrated approach. AWS CodeBuild is fully integrated with AWS CodePipeline and can be used to run unit tests as part of the CI/CD process. Placing the testing stage before deployment ensures that only tested code is deployed. The buildspec can be configured to fail the build if tests do not pass, and CodeBuild's test reports feature allows for easy viewing and analysis of test results.

upvoted 6 times

65703c1 Most Recent 8 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

xxxx1 11 months ago

c is the correct answer

upvoted 1 times

NinjaCloud 1 year, 3 months ago

Correct answer: B

upvoted 1 times

Gold07 1 year, 3 months ago

c is the correct answer

upvoted 1 times

Cerakoted 1 year, 3 months ago

Selected Answer: C

I think C is correct.

Typical consists of stages are..

Build -> Test -> Deploy(test) -> Load Test -> and others

upvoted 2 times

dilleman 1 year, 3 months ago

Selected Answer: C

C should be correct.

upvoted 3 times



  **Digo30sp** 1 year, 3 months ago

Selected Answer: B

The correct answer is (B).

Solution (B) is the simplest and requires the least operational effort. It involves adding a new stage to the CodePipeline pipeline that uses AWS CodeBuild to run the unit tests. The CodeBuild stage can be configured to fail if any tests fail. The CodeBuild test report can be integrated into the CodeBuild console so that developers can view test results.

upvoted 1 times

  **04075e0** 2 months, 4 weeks ago



definitely not B, since nobody doing tests after deployment. No sense

upvoted 1 times

  **dilleman** 1 year, 3 months ago

This does not make sense. Why run the tests after the deploy when you can choose option C, to run the tests before the deploy? C should be best practice and the same amount of effort as B.

upvoted 5 times

  **Dibaal** 1 year, 3 months ago

funny 😄

upvoted 1 times

A company has multiple Amazon VPC endpoints in the same VPC. A developer needs to configure an Amazon S3 bucket policy so users can access an S3 bucket only by using these VPC endpoints.

Which solution will meet these requirements?

- A. Create multiple S3 bucket policies by using each VPC endpoint ID that have the aws:SourceVpce value in the StringNotEquals condition.
- B. Create a single S3 bucket policy that has the aws:SourceVpc value and in the StringNotEquals condition to use VPC ID.
- C. Create a single S3 bucket policy that has the aws:SourceVpce value and in the StringNotEquals condition to use vpce*.
- D. Create a single S3 bucket policy that has multiple aws:sourceVpce value in the StringNotEquals condition. Repeat for all the VPC endpoint IDs.

Correct Answer: D

Community vote distribution

D (86%)

10%

 **dilleman** Highly Voted 1 year, 3 months ago

Selected Answer: D

C works as well but It is a broad solution I think it's better practice to use D and specify the exact endpoints that the user can access from.

```
"aws:sourceVpce": ["vpce-id1", "vpce-id2", "..."]
```

upvoted 7 times

 **CrescentShared** Highly Voted 1 year, 2 months ago

I don't think any of the options is correct. Seriously StringNotEquals not StringEquals?

upvoted 6 times

 **shake76** 1 year, 1 month ago

I think the same "A developer needs to configure an Amazon S3 bucket policy so users can access an S3 bucket only by using these VPC endpoints"

upvoted 1 times

 **vipyodha** 10 months, 3 weeks ago

StringNotEqual is for the deny of other that mentioned vpce.

```
{
  "Version": "2012-10-17",
  "Id": "Policy1415115909152",
  "Statement": [
    {
      "Sid": "Access-to-specific-VPCE-only",
      "Principal": "*",
      "Action": "s3:*",
      "Effect": "Deny",
      "Resource": ["arn:aws:s3:::awsexamplebucket1",
        "arn:aws:s3:::awsexamplebucket1/*"],
      "Condition": {
        "StringNotEquals": {
          "aws:SourceVpce": "vpce-1a2b3c4d"
        }
      }
    }
  ]
}
```

upvoted 2 times

 **examtopics111** 6 months ago


for bucket policy, if vpce isnt explicitly allowed, it's by default denied anyway so it should have been allow string equal vpce?
upvoted 1 times

 **albert_kuo** Most Recent 4 months, 1 week ago

Selected Answer: D

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyAccessIfNotFromSpecificVPCEndpoints",
      "Effect": "Deny",
      "Principal": "*",
      "Action": "s3:*",
      "Resource": [
        "arn:aws:s3:::your-bucket-name",
        "arn:aws:s3:::your-bucket-name/*"
      ],
      "Condition": {
        "StringNotEquals": {
          "aws:SourceVpce": [
            "vpce-0123456789abcdef0",
            "vpce-0fedcba9876543210",
            "vpce-0a1b2c3d4e5f6a7b8"
          ]
        }
      }
    }
  ]
}
```

upvoted 3 times

 **65703c1** 8 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

 **KarBiswa** 11 months, 1 week ago

Selected Answer: D

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-bucket-policies-vpc-endpoint.html> typically explained the same scenario. D beyond doubt.

upvoted 2 times

 **cajilaxu** 12 months ago

Selected Answer: D

D is right answer

Get up-to-date <https://www.pinterest.com/pin/937522847419120392>

upvoted 1 times

 **joshnort** 12 months ago

Selected Answer: D

D, based on the following documentation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-bucket-policies-vpc-endpoint.html#example-bucket-policies-restrict-accesss-vpc-endpoint>

upvoted 1 times

 **CrescentShared** 1 year ago

Why it's StringNotEquals instead of StringEquals? Is the question wrong or my English is too bad to understand this?

upvoted 2 times

 **vipyodha** 10 months, 3 weeks ago

It is StringNotEqual, means if source vpce is not this then deny access

```
{
  "Version": "2012-10-17",
  "Id": "Policy1415115909152",
  "Statement": [
    {
      "Sid": "Access-to-specific-VPCE-only",
      "Principal": "*",
      "Action": "s3:*",
      "Effect": "Deny",
      "Resource": ["arn:aws:s3:::awsexamplebucket1",
        "arn:aws:s3:::awsexamplebucket1/*"],
      "Condition": {
        "StringNotEquals": {
          "aws:SourceVpce": "vpce-1a2b3c4d"
        }
      }
    }
  ]
}
```

upvoted 3 times

SerialiDr 1 year ago

Selected Answer: D

This option is the closest to being correct, but it should use StringEquals instead of StringNotEquals. The correct approach is to use a single S3 bucket policy with a condition that includes aws:SourceVpce with StringEquals for the specific VPC endpoint IDs. This will ensure that access is allowed only from those specified endpoints.

upvoted 2 times

rrshah83 1 year, 1 month ago

Selected Answer: B

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-bucket-policies-vpc-endpoint.html#example-bucket-policies-restrict-access-vpc>

upvoted 2 times

Certified101 1 year, 1 month ago

Selected Answer: D

D is correct

upvoted 1 times

PrakashM14 1 year, 3 months ago

Selected Answer: D

in option C :

Condition": {

"StringNotEqualsIfExists": {

"aws:sourceVpce": "vpce*",

}

}

it might Deny access from all VPC endpoints.

so the ans is D

upvoted 2 times

ekutas 1 year, 2 months ago



D says "aws:sourceVpce value in the StringNotEquals condition". StringNotEquals won't work, it deny access for specified VPC ids

upvoted 1 times

ekutas 1 year, 2 months ago

Of course if we use "Effect": "Allow"))

upvoted 2 times

  **Digo30sp** 1 year, 3 months ago

Selected Answer: C

The correct answer is (C).

Solution (C) is the simplest and will meet the company's requirements. It creates a single S3 bucket policy that has the value `aws:SourceVpce` and the `StringNotEquals` condition to use `vpce*`. This will only allow users who are using a VPC endpoint in the same VPC to access the S3 bucket.

upvoted 1 times

A company uses a custom root certificate authority certificate chain (Root CA Cert) that is 10 KB in size to generate SSL certificates for its on-premises HTTPS endpoints. One of the company's cloud-based applications has hundreds of AWS Lambda functions that pull data from these endpoints. A developer updated the trust store of the Lambda execution environment to use the Root CA Cert when the Lambda execution environment is initialized. The developer bundled the Root CA Cert as a text file in the Lambda deployment bundle.

After 3 months of development, the Root CA Cert is no longer valid and must be updated. The developer needs a more efficient solution to update the Root CA Cert for all deployed Lambda functions. The solution must not include rebuilding or updating all Lambda functions that use the Root CA Cert. The solution must also work for all development, testing, and production environments. Each environment is managed in a separate AWS account.

Which combination of steps should the developer take to meet these requirements MOST cost-effectively? (Choose two.)

- A. Store the Root CA Cert as a secret in AWS Secrets Manager. Create a resource-based policy. Add IAM users to allow access to the secret.
- B. Store the Root CA Cert as a SecureString parameter in AWS Systems Manager Parameter Store. Create a resource-based policy. Add IAM users to allow access to the policy.
- C. Store the Root CA Cert in an Amazon S3 bucket. Create a resource-based policy to allow access to the bucket.
- D. Refactor the Lambda code to load the Root CA Cert from the Root CA Cert's location. Modify the runtime trust store inside the Lambda function handler.
- E. Refactor the Lambda code to load the Root CA Cert from the Root CA Cert's location. Modify the runtime trust store outside the Lambda function handler.

Correct Answer: AE

Community vote distribution



kiwtirApp Highly Voted 1 year, 3 months ago

Selected Answer: AE

The max size of storage in Secrets Manager is 10kb. For SSM Parameter store, it's 8Kb.

Correct options are A and E.

upvoted 11 times

not_a_bot_definitely 1 year, 2 months ago

Secrets Manager is not cost-effective compared to option C - S3 bucket.

Question clearly asks "MOST cost-effective"

<https://www.examttopics.com/discussions/amazon/view/96242-exam-aws-certified-developer-associate-topic-1-question-429/>

So answer is CE

upvoted 7 times

ShakthiGCP Most Recent 2 months, 3 weeks ago

Selected Answer: AE

beyond cost effective. AWS always recommend the secured way. in that case, it will be A and E

upvoted 1 times

9d8dd9c 3 months ago

In both A and B why does it say IAM user? should it not be IAM role? IAM role for Lambda to access them?

upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: AE

AE is the correct answer.

upvoted 1 times

🗳️ 👤 **Melisa202401** 9 months, 4 weeks ago

Selected Answer: BE

This solution will meet the requirements by storing the Root CA Cert as a Secure String parameter in AWS Systems Manager Parameter Store, which is a secure and scalable service for storing and managing configuration data and secrets. The resource-based policy will allow IAM users in different AWS accounts and environments to access the parameter without requiring cross-account roles or permissions. The Lambda code will be refactored to load the Root CA Cert from the parameter store and modify the runtime trust store outside the Lambda function handler, which will improve performance and reduce latency by avoiding repeated calls to Parameter Store and trust store modifications for each invocation of the Lambda function.

upvoted 2 times

🗳️ 👤 **yingying920928** 10 months, 2 weeks ago

Selected Answer: CE

Cost effective, use S3 instead of Secrets Manager.

upvoted 2 times

🗳️ 👤 **KarBiswa** 11 months, 1 week ago

Selected Answer: AE

After going through the links :

A : <https://aws.amazon.com/blogs/security/use-aws-secrets-manager-to-simplify-the-management-of-private-certificates/>

E : <https://docs.aws.amazon.com/acm/latest/userguide/renew-private-cert.html>

upvoted 1 times

🗳️ 👤 **dostonbekabdullaev** 1 year ago

Selected Answer: CE

C.E.

Secrets Manager is the most expensive amongst all options. S3 seems more cost-effective.

B. is incorrect, because at the end it states about accessing to the policy, not to the parameter itself.

upvoted 3 times

🗳️ 👤 **SerialiDr** 1 year ago

Selected Answer: BE

Also AE works, but BE is more cost effective.

upvoted 1 times

🗳️ 👤 **dostonbekabdullaev** 1 year ago

Pay attention on this part "Add IAM users to allow access to the policy." It should give an access to the parameter, not to the policy.

upvoted 3 times

🗳️ 👤 **CalvinL4** 1 year ago

CE should be the answer. The string size is over 4/8 kb which the parameter store allows. So, the parameter store is out. Comparing the price, s3 is much cheaper than secrets manager.

upvoted 2 times

🗳️ 👤 **rrshah83** 1 year, 1 month ago

Selected Answer: AE

<https://aws.amazon.com/blogs/security/use-aws-secrets-manager-to-simplify-the-management-of-private-certificates/>

upvoted 1 times

🗳️ 👤 **rrshah83** 1 year, 1 month ago

Selected Answer: AE

can you do resource based policies for param store?

upvoted 1 times

🗳️ 👤 **Hanny** 1 year, 1 month ago

Selected Answer: CE

<https://www.examttopics.com/discussions/amazon/view/96242-exam-aws-certified-developer-associate-topic-1-question-429/>

upvoted 1 times

🗳️ 👤 **tqiu654** 1 year, 1 month ago

Selected Answer: BD

CHatGPT: BD

upvoted 1 times

🗳️ 👤 **wonder_man** 1 year, 3 months ago

Selected Answer: CE

I can't see why using AWS Secrets Manager can be cost-effective, so I'm voting for C
upvoted 4 times

  **Rameez1** 1 year, 3 months ago

Selected Answer: BE

Using Parameter store is more cost effective then secrets manager.
upvoted 2 times

  **TallManDan** 1 year, 3 months ago

Secrets Manager is an additional cost over Parameter Store. So if you see a question that looks for the least amount of overhead, Secrets Manager is much more versatile. But for least amount of cost, Parameter Store is included with the service for no additional costs.
upvoted 3 times

A developer maintains applications that store several secrets in AWS Secrets Manager. The applications use secrets that have changed over time. The developer needs to identify required secrets that are still in use. The developer does not want to cause any application downtime.

What should the developer do to meet these requirements?

- A. Configure an AWS CloudTrail log file delivery to an Amazon S3 bucket. Create an Amazon CloudWatch alarm for the GetSecretValue Secrets Manager API operation requests.
- B. Create a secretsmanager-secret-unused AWS Config managed rule. Create an Amazon EventBridge rule to initiate notifications when the AWS Config managed rule is met.
- C. Deactivate the applications secrets and monitor the applications error logs temporarily.
- D. Configure AWS X-Ray for the applications. Create a sampling rule to match the GetSecretValue Secrets Manager API operation requests.

Correct Answer: B

Community vote distribution

B (71%)

A (29%)

 **chris_777** Highly Voted 1 year, 2 months ago

Selected Answer: B


I think B is correct <https://docs.aws.amazon.com/config/latest/developerguide/secretsmanager-secret-unused.html>

- A. could work but requires additional work to identify unused secrets.
 - C. is too risky and could cause downtime.
 - D. not the right use case
- upvoted 7 times

 **9d8dd9c** Most Recent 3 months ago


With option A, if you want to know what secrets are not used anymore for the past 90 days, then you need to wait 90 days, get a list of the ones that are used then minus this from the total list... but option B returns the unused list for the last 90 days directly

upvoted 1 times

 **MasoudK** 3 months, 4 weeks ago

source copilot: Why Option A is Better:

1. Real-Time Monitoring: By using CloudTrail and CloudWatch, you can monitor GetSecretValue API calls in real-time, providing immediate insights into which secrets are being accessed.
 2. Detailed Logging: CloudTrail logs provide detailed information about each API call, including the source IP, user, and timestamp, which can help identify the specific applications or instances accessing the secrets.
 3. Customizable Alerts: CloudWatch alarms can be configured to trigger notifications based on specific patterns or thresholds, offering more flexibility and control over monitoring.
- upvoted 1 times

 **MasoudK** 3 months, 4 weeks ago

Why Option B Might Not Be the Best Choice:

1. Rule Limitation: The secretsmanager-secret-unused AWS Config managed rule checks if a secret has not been retrieved for a specified period. However, it might not provide real-time insights into which secrets are currently in use.
 2. Delayed Detection: This rule might only detect secrets that have not been used for a while, potentially missing secrets that are infrequently accessed but still required.
 3. Lack of Granularity: The rule might not provide detailed information about the specific applications or instances accessing the secrets, making it harder to pinpoint which secrets are actively used.
- upvoted 1 times

 **65703c1** 8 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

 **KarBiswa** 11 months, 1 week ago

Selected Answer: B

We need "secrets that are still in use". "B" secretsmanager-secret-unused returns unused. So we can easily determine the used secrets if it is not falling under this scanner

upvoted 2 times

🗨️ **rimaSamir** 1 year ago

My choice is "A".

We need "secrets that are still in use". "B" secretsmanager-secret-unused returns unused.

upvoted 1 times

🗨️ **SerialIDr** 1 year ago

Selected Answer: A

A. Use an AWS Step Functions state machine to monitor API failures. Use the Wait state to delay calling the Lambda function: This is a viable and efficient solution. AWS Step Functions can orchestrate the Lambda function invocations and manage the workflow, including handling API call rate limits. The Wait state can be used to introduce delays between API calls to ensure compliance with the rate limits. This approach also allows for handling errors and retries effectively.

B. Use an Amazon Simple Queue Service (Amazon SQS) queue to hold the API calls. Configure the Lambda function to poll the queue within the API threshold limits: While using SQS to queue API call requests is a good way to manage workload, it adds complexity to the solution. The Lambda function would need to be modified to manage the queue and ensure API calls are made within the threshold limits. This approach might not be as straightforward and efficient as using Step Functions.

upvoted 2 times

🗨️ **tqiu654** 1 year, 1 month ago

Selected Answer: A

ChatGPT:A

upvoted 1 times

🗨️ **kaes** 1 year, 2 months ago

It's easier to use a built-in solution in AWS Config (check chris_777 answer)

upvoted 1 times

🗨️ **kaes** 1 year, 2 months ago

Selected Answer: B

It's easier to use a default built-in solution in AWS Config (check chris_777 answer)

upvoted 1 times

🗨️ **CrescentShared** 1 year, 2 months ago

Selected Answer: A

I think A is a more direct way, while B needs an inference after receiving the notification for 'unused'.

upvoted 1 times

🗨️ **LemonGremlin** 1 year, 3 months ago

Selected Answer: B

B is correct for this one.

upvoted 1 times

🗨️ **dilleman** 1 year, 3 months ago

Selected Answer: A

A is correct. . AWS CloudTrail can track API calls, including the GetSecretValue call for AWS Secrets Manager. By setting up CloudTrail log delivery to an S3 bucket, the developer can analyze which secrets are being accessed. Using CloudWatch to create an alarm for the GetSecretValue API call provides insight into which secrets are actively being retrieved, thus indicating which secrets are in use.

upvoted 2 times

🗨️ **dilleman** 1 year, 3 months ago

I think i change my mind to B. B Must be correct..

upvoted 3 times

🗨️ **CrescentShared** 1 year, 2 months ago

Why did you change your mind, please? A looks super correct to me.

upvoted 1 times

🗨️ **Digo30sp** 1 year, 3 months ago

Selected Answer: B

The correct answer is (B).

Solution (B) is the best option to meet the developer's requirements. It allows the developer to identify necessary secrets that are still in use without causing any application downtime.

upvoted 3 times

A developer is writing a serverless application that requires an AWS Lambda function to be invoked every 10 minutes.

What is an automated and serverless way to invoke the function?

- A. Deploy an Amazon EC2 instance based on Linux, and edit its `/etc/crontab` file by adding a command to periodically invoke the Lambda function.
- B. Configure an environment variable named `PERIOD` for the Lambda function. Set the value to 600.
- C. Create an Amazon EventBridge rule that runs on a regular schedule to invoke the Lambda function.
- D. Create an Amazon Simple Notification Service (Amazon SNS) topic that has a subscription to the Lambda function with a 600-second timer.

Correct Answer: C

Community vote distribution

C (100%)

🗉 👤 **tsangckl** 7 months, 2 weeks ago

This appear at 17 Jun exam
upvoted 1 times

🗉 👤 **65703c1** 8 months, 1 week ago

Selected Answer: C

C is the correct answer.
upvoted 2 times

🗉 👤 **SerialiDr** 1 year ago

Selected Answer: C

C. Create an Amazon EventBridge rule that runs on a regular schedule to invoke the Lambda function: This is the correct and most suitable option. Amazon EventBridge (formerly CloudWatch Events) allows you to set up rules that trigger on a schedule. You can create a rule with a cron or rate expression to invoke the Lambda function every 10 minutes. This approach is fully serverless and does not require managing any servers or additional infrastructure.

upvoted 3 times

🗉 👤 **dilleman** 1 year, 3 months ago

Selected Answer: C

C is correct.
Amazon EventBridge can be used to run Lambda functions on a regular schedule. You can set a cron or rate expression to define the schedule.
upvoted 3 times

🗉 👤 **Digo30sp** 1 year, 3 months ago

Selected Answer: C

The correct answer is (C).

Solution (C) is the best option to meet the developer's requirements. It allows the developer to invoke the Lambda function in an automated and serverless way.

upvoted 2 times

A company is using Amazon OpenSearch Service to implement an audit monitoring system. A developer needs to create an AWS CloudFormation custom resource that is associated with an AWS Lambda function to configure the OpenSearch Service domain. The Lambda function must access the OpenSearch Service domain by using OpenSearch Service internal master user credentials.

What is the MOST secure way to pass these credentials to the Lambda function?

- A. Use a CloudFormation parameter to pass the master user credentials at deployment to the OpenSearch Service domain's MasterUserOptions and the Lambda function's environment variable. Set the NoEcho attribute to true.
- B. Use a CloudFormation parameter to pass the master user credentials at deployment to the OpenSearch Service domain's MasterUserOptions and to create a parameter in AWS Systems Manager Parameter Store. Set the NoEcho attribute to true. Create an IAM role that has the ssm:GetParameter permission. Assign the role to the Lambda function. Store the parameter name as the Lambda function's environment variable. Resolve the parameter's value at runtime.
- C. Use a CloudFormation parameter to pass the master user credentials at deployment to the OpenSearch Service domain's MasterUserOptions and the Lambda function's environment variable. Encrypt the parameter's value by using the AWS Key Management Service (AWS KMS) encrypt command.
- D. Use CloudFormation to create an AWS Secrets Manager secret. Use a CloudFormation dynamic reference to retrieve the secret's value for the OpenSearch Service domain's MasterUserOptions. Create an IAM role that has the secretsmanager:GetSecretValue permission. Assign the role to the Lambda function. Store the secret's name as the Lambda function's environment variable. Resolve the secret's value at runtime.

Correct Answer: D

Community vote distribution

D (100%)

65703c1 2 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 2 times

SerialiDr 6 months, 2 weeks ago

Selected Answer: D

This approach is the most secure and aligns with best practices for managing secrets. The credentials are stored in AWS Secrets Manager, which is specifically designed for managing and protecting secrets. The credentials are retrieved dynamically at runtime by the Lambda function, and the use of IAM roles ensures that only the Lambda function has access to these secrets. This method also benefits from the security and rotation features of AWS Secrets Manager.

upvoted 2 times

dilleman 9 months, 3 weeks ago

Selected Answer: D

D is correct.

upvoted 2 times

Digo30sp 10 months ago

Selected Answer: D

The correct answer is (D).

Solution (D) is the most secure way to pass the credentials to the Lambda function because it uses AWS Secrets Manager to store the credentials in encrypted form.

upvoted 3 times

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 ElastiCache.
- B. Write data to Amazon Elastic Block Store.
- C. Write data to Amazon EC2 Instance Store.
- D. Write data to the root filesystem.

Correct Answer: A

Community vote distribution

A (100%)

🗨️ **65703c1** 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 2 times

🗨️ **SerialiDr** 6 months, 2 weeks ago

Selected Answer: A

Amazon ElastiCache provides a fast, in-memory data store or cache. It is often used for session management in distributed applications. Data stored in ElastiCache can be accessed quickly and reliably by any of the EC2 instances behind the ELB, making it an ideal choice for session data that needs to be shared across multiple servers.

upvoted 2 times

🗨️ **dilleman** 9 months, 3 weeks ago

Selected Answer: A

A is correct.

By storing session data in ElastiCache, you ensure that regardless of which EC2 instance handles a given request, the session data can be consistently and rapidly accessed.

upvoted 3 times

🗨️ **Digo30sp** 10 months ago

Selected Answer: A

The correct answer is (A).

Amazon ElastiCache is a distributed memory caching solution that is ideal for session data. ElastiCache provides high-performance and durable session data storage that can be shared across multiple EC2 instances.

upvoted 3 times

An ecommerce application is running behind an Application Load Balancer. A developer observes some unexpected load on the application during non-peak hours. The developer wants to analyze patterns for the client IP addresses that use the application.

Which HTTP header should the developer use for this analysis?

- A. The X-Forwarded-Proto header
- B. The X-Forwarded-Host header
- C. The X-Forwarded-For header
- D. The X-Forwarded-Port header

Correct Answer: C

Community vote distribution

C (100%)

 **chris_777** Highly Voted 9 months ago

Selected Answer: C

C is correct.

X-Forwarded-Proto: protocol (HTTP/HTTPS)

X-Forwarded-Host: original Host header requested by the client

X-Forwarded-For: original IP address of a client (CORRECT)

X-Forwarded-Port header: original port that the client used to connect

upvoted 9 times

 **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: C

C is the correct answer.


upvoted 1 times

 **tapan666** 9 months, 1 week ago

Selected Answer: C

C is correct


upvoted 1 times

 **dilleman** 9 months, 3 weeks ago

Selected Answer: C

C is correct

upvoted 1 times

 **Cerakoted** 9 months, 3 weeks ago

Selected Answer: C

X-Forwarded-For HTTP header contains the IP address of the original client

upvoted 1 times

 **Digo30sp** 10 months ago

Selected Answer: C

The correct answer is (C).

The X-Forwarded-For HTTP header contains the IP address of the original client that made the request. The developer can use this header to analyze patterns for the IP addresses of clients using the application.

upvoted 2 times

A developer migrated a legacy application to an AWS Lambda function. The function uses a third-party service to pull data with a series of API calls at the end of each month. The function then processes the data to generate the monthly reports. The function has been working with no issues so far.


The third-party service recently issued a restriction to allow a fixed number of API calls each minute and each day. If the API calls exceed the limit for each minute or each day, then the service will produce errors. The API also provides the minute limit and daily limit in the response header. This restriction might extend the overall process to multiple days because the process is consuming more API calls than the available limit.

What is the MOST operationally efficient way to refactor the serverless application to accommodate this change?

- A. Use an AWS Step Functions state machine to monitor API failures. Use the Wait state to delay calling the Lambda function.
- B. Use an Amazon Simple Queue Service (Amazon SQS) queue to hold the API calls. Configure the Lambda function to poll the queue within the API threshold limits.
- C. Use an Amazon CloudWatch Logs metric to count the number of API calls. Configure an Amazon CloudWatch alarm that stops the currently running instance of the Lambda function when the metric exceeds the API threshold limits.
- D. Use Amazon Kinesis Data Firehose to batch the API calls and deliver them to an Amazon S3 bucket with an event notification to invoke the Lambda function.

Correct Answer: A

Community vote distribution



wonder_man Highly Voted 1 year, 3 months ago

Selected Answer: A

B: I don't see how the Lamba function can be configured this way
upvoted 6 times

trieudo Most Recent 1 month, 1 week ago

Selected Answer: A

A: makes the most sense in this question: Use an AWS Step Functions state machine to monitor API failures. Use the Wait state to delay calling the Lambda function.

B: SQS queue wouldn't work because the maximum delay allowed in an SQS queue is 15 minutes. ==> 'extend the overall process to multiple days' will be violated
upvoted 1 times

trieudo 1 month, 1 week ago

refer: https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-delay-queues.html?utm_source=chatgpt.com
upvoted 1 times

tsangckl 7 months, 2 weeks ago

This appear at 17 Jun exam
upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: A

A is the correct answer.
upvoted 1 times

KarBiswa 11 months, 1 week ago

Selected Answer: C

I would go for option C because response value always contains the limit value and which can be retrieved into the cloudwatch and can be used to block the the lambda calls
upvoted 1 times

🗨️ 👤 **KarBiswa** 10 months, 3 weeks ago

Changing to A

upvoted 3 times

🗨️ 👤 **KillThemWithKindness** 11 months, 1 week ago

Selected Answer: A

The solution that will meet the requirements is to use an AWS Step Functions state machine to monitor API failures. Use the Wait state to delay calling the Lambda function. This way, the developer can refactor the serverless application to accommodate the change in a way that is automated and scalable. The developer can use Step Functions to orchestrate the Lambda function and handle any errors or retries. The developer can also use the Wait state to pause the execution for a specified duration or until a specified timestamp, which can help avoid exceeding the API limits. The other options either involve using additional services that are not necessary or appropriate for this scenario, or do not address the issue of API failures.

<https://www.freecram.net/question/Amazon.DVA-C02.v2023-11-29.q68/a-developer-migrated-a-legacy-application-to-an-aws-lambda-function-the-function-uses-a-third-party-18#>

upvoted 3 times

🗨️ 👤 **konieczny69** 12 months ago

Selected Answer: A

Who is going to orchestrate lambda invocation?

SQS is for decoupling, not for scheduled invocations.

A is the only option.

upvoted 4 times

🗨️ 👤 **CrescentShared** 1 year ago

Selected Answer: B

Option A with AWS Step Functions can handle the frequency of API calls by introducing a delay (Wait state) between retries after a failure due to rate limiting, it doesn't inherently solve the problem of the total number of calls per day. If the total number of necessary API calls exceeds the daily limit set by the third-party service, simply adding a delay between retries will not prevent the overall daily limit from being exceeded.

upvoted 3 times

🗨️ 👤 **SerialiDr** 1 year ago

Selected Answer: A

A. Use an AWS Step Functions state machine to monitor API failures. Use the Wait state to delay calling the Lambda function: This is a viable and efficient solution. AWS Step Functions can orchestrate the Lambda function invocations and manage the workflow, including handling API call rate limits. The Wait state can be used to introduce delays between API calls to ensure compliance with the rate limits. This approach also allows for handling errors and retries effectively.

B. Use an Amazon Simple Queue Service (Amazon SQS) queue to hold the API calls. Configure the Lambda function to poll the queue within the API threshold limits: While using SQS to queue API call requests is a good way to manage workload, it adds complexity to the solution. The Lambda function would need to be modified to manage the queue and ensure API calls are made within the threshold limits. This approach might not be as straightforward and efficient as using Step Functions.

upvoted 3 times

🗨️ 👤 **JohnPI** 1 year ago

Selected Answer: B

B is the most operationally efficient way

upvoted 1 times

🗨️ 👤 **Snape** 1 year ago

Selected Answer: B

b is the answer

upvoted 1 times

🗨️ 👤 **rrshah83** 1 year, 1 month ago

Selected Answer: B

sqz decouples lambda from api service

upvoted 1 times

🗨️ 👤 **chewasa** 1 year, 1 month ago

Selected Answer: B

While Step Functions can be used for workflow orchestration, it may not be the most straightforward solution for handling rate limits in this scenario.

upvoted 1 times

🗨️ 👤 **tqiu654** 1 year, 1 month ago

Selected Answer: A

ChatGPT: A

upvoted 2 times

🗨️ 👤 **ShawnWon** 1 year, 2 months ago

B.

Option A (AWS Step Functions) might introduce unnecessary complexity and does not directly address the need to control the rate of API calls within the specified limits.

Option C (CloudWatch Logs metric and alarm) provides monitoring capabilities but doesn't offer a direct mechanism to control the rate of API calls within the Lambda function.

Option D (Kinesis Data Firehose) is designed for real-time streaming and might not be the most suitable option for this scenario, as it may not provide the fine-grained control needed to stay within the API call limits.

upvoted 3 times

🗨️ 👤 **dilleman** 1 year, 3 months ago

Selected Answer: A

A is Correct. AWS Step Functions can be used to create a workflow to handle the API calls. You can make the Lambda function inspect the response headers from the third-party service to determine the current API call limits and then pass that to the Wait state of the state machine for proper delays.

upvoted 2 times

🗨️ 👤 **Digo30sp** 1 year, 3 months ago

Selected Answer: B

The correct answer is (B).

Solution (B) is the most operationally efficient way to refactor the serverless application to accommodate this change. This solution allows the Lambda function to continue executing API calls even if the API call limit is reached. The Amazon SQS queue will act as a buffer for API calls that exceed the limit. The Lambda function can then poll the queue within the API limits.

upvoted 4 times

A developer must analyze performance issues with production-distributed applications written as AWS Lambda functions. These distributed Lambda applications invoke other components that make up the applications.

How should the developer identify and troubleshoot the root cause of the performance issues in production?

- A. Add logging statements to the Lambda functions, then use Amazon CloudWatch to view the logs.
- B. Use AWS CloudTrail and then examine the logs.
- C. Use AWS X-Ray, then examine the segments and errors.
- D. Run Amazon Inspector agents and then analyze performance.

Correct Answer: C

Community vote distribution

C (100%)

 **Digo30sp** Highly Voted 10 months ago

Selected Answer: C

The correct answer is (C).

AWS X-Ray is the best tool for identifying and addressing the root cause of performance issues in distributed production applications. X-Ray provides an overview of the entire call stack, including the Lambda functions and other components they invoke.


upvoted 5 times

 **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

 **dilleman** 9 months, 3 weeks ago

Selected Answer: C

C is correct.

upvoted 4 times

A developer wants to deploy a new version of an AWS Elastic Beanstalk application. During deployment, the application must maintain full capacity and avoid service interruption. Additionally, the developer must minimize the cost of additional resources that support the deployment.

Which deployment method should the developer use to meet these requirements?

- A. All at once
- B. Rolling with additional batch
- C. Blue/green
- D. Immutable

Correct Answer: B

Community vote distribution



🗨️ **Nagorasacle** Highly Voted 1 year, 3 months ago

Selected Answer: B

B: Rolling with additional batch , considering "minimize the cost of additional resources"
C costly than B, due to double capacity
upvoted 15 times

🗨️ **SerialIDr** Highly Voted 1 year ago

Selected Answer: D

D. Immutable

The immutable method strikes a balance between maintaining service availability and controlling costs. It avoids the downtime associated with the all-at-once method and doesn't require the more extensive resource duplication of the blue/green method. While it does temporarily increase resource usage (similar to rolling with an additional batch), it's generally more efficient and less risky than updating instances in-place.

upvoted 5 times

🗨️ **aws_god** 3 months ago

you mean D

upvoted 1 times

🗨️ **ShakthiGCP** Most Recent 2 months, 3 weeks ago

Selected Answer: B

Rolling deployments reuse existing instances for a portion of the deployment time. While new instances are launched, old ones continue to serve traffic. This minimizes idle compute time and reduces costs.

upvoted 1 times

🗨️ **65703c1** 8 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 2 times

🗨️ **a5fc516** 9 months, 3 weeks ago

Ans-B: The immutable deployment method updates an application by launching new instances with the new version in a new Auto Scaling group, alongside the existing instances running the old version. This ensures that the application maintains full capacity during the deployment because the existing environment is unaffected until the new environment is fully deployed and verified. Once the deployment is successful, traffic is shifted to the new instances, and the old ones are terminated. This minimizes downtime and provides a quick rollback if needed. The cost of additional resources is limited to the duration of the deployment, after which the old resources are terminated.

upvoted 1 times

🗨️ **aws_god** 3 months ago

you mean D

upvoted 1 times

🗨️ **KarBiswa** 11 months, 1 week ago

Selected Answer: B

[https://docs.aws.amazon.com/whitepapers/latest/practicing-continuous-integration-continuous-delivery/deployment-](https://docs.aws.amazon.com/whitepapers/latest/practicing-continuous-integration-continuous-delivery/deployment-methods.html#:~:text=During%20the%20deployment%20process%20two%20software%20versions%2C%20new%20and%20old%2C%20are%20running%20o)

[methods.html#:~:text=During%20the%20deployment%20process%20two%20software%20versions%2C%20new%20and%20old%2C%20are%20running%20o](https://docs.aws.amazon.com/whitepapers/latest/practicing-continuous-integration-continuous-delivery/deployment-methods.html#:~:text=During%20the%20deployment%20process%20two%20software%20versions%2C%20new%20and%20old%2C%20are%20running%20o)

It does not need a new instance

upvoted 1 times

🗨️ **Roimasu** 1 year, 3 months ago

Selected Answer: D

This method performs updates by launching a new set of instances in a new Auto Scaling group. Once the new instances pass health checks, they are moved into the existing Auto Scaling group, and the old instances are terminated. This method ensures full capacity, avoids downtime, and minimizes additional costs because it does not double the environment's running resources for an extended period. It adds resources temporarily and only in the amount necessary to maintain capacity.

upvoted 3 times

🗨️ **NinjaCloud** 1 year, 3 months ago

Should be B "Ultimately, the choice between "Rolling with additional batch" and "Blue/green" deployments should depend on your specific requirements and constraints. If maintaining full capacity is a crucial factor, then "Rolling with additional batch" could be the better choice."

upvoted 4 times

🗨️ **ut18** 1 year, 3 months ago

MS Bing answer: B vs Chag GPT answer: C

Your choice?

upvoted 1 times

🗨️ **CrescentShared** 1 year, 2 months ago

ChatGPT4 changed its mind to select D today.

upvoted 1 times

🗨️ **Learning4life** 1 year, 3 months ago

C and D are wrong, since they both require additional resources.

upvoted 1 times

🗨️ **joosh96** 1 year, 3 months ago

Selected Answer: C

chat gpt replied

upvoted 1 times

🗨️ **Cerakoted** 1 year, 3 months ago

Selected Answer: B

Answer is B

One of requirement - the developer [must minimize the cost of additional resources] that support the deployment.

upvoted 4 times

🗨️ **dilleman** 1 year, 3 months ago

Selected Answer: D

I vote for D since the requirement is to minimize the costs of resources. Blue/green is a good and safe way to solve this but it costs more resources than an Immutable rollout.

Immutable: Launches a new set of instances in a new temporary environment to ensure that the new version works as expected. Once the new version is verified, traffic is rerouted to the new set of instances, and the old instances are terminated. This method maintains full capacity, avoids service interruptions, and minimizes the cost compared to blue/green deployments since the overlap in running resources is shorter.

upvoted 4 times

🗨️ **Digo30sp** 1 year, 3 months ago

Selected Answer: C

The correct answer is (C).

The blue/green deployment method is the best option to meet the developer's requirements. Blue/green allows the developer to deploy a new version of the application without service interruption. This is done by creating a blue production environment and a green production environment. The blue environment is the current production environment and the green environment is the new version of the application. The developer can then test the new version of the application in the green environment before putting it into production.

upvoted 2 times

A developer has observed an increase in bugs in the AWS Lambda functions that a development team has deployed in its Node.js application. To minimize these bugs, the developer wants to implement automated testing of Lambda functions in an environment that closely simulates the Lambda environment.

The developer needs to give other developers the ability to run the tests locally. The developer also needs to integrate the tests into the team's continuous integration and continuous delivery (CI/CD) pipeline before the AWS Cloud Development Kit (AWS CDK) deployment.

Which solution will meet these requirements?

- A. Create sample events based on the Lambda documentation. Create automated test scripts that use the `cdk local invoke` command to invoke the Lambda functions. Check the response. Document the test scripts for the other developers on the team. Update the CI/CD pipeline to run the test scripts.
- B. Install a unit testing framework that reproduces the Lambda execution environment. Create sample events based on the Lambda documentation. Invoke the handler function by using a unit testing framework. Check the response. Document how to run the unit testing framework for the other developers on the team. Update the CI/CD pipeline to run the unit testing framework.
- C. Install the AWS Serverless Application Model (AWS SAM) CLI tool. Use the `sam local generate-event` command to generate sample events for the automated tests. Create automated test scripts that use the `sam local invoke` command to invoke the Lambda functions. Check the response. Document the test scripts for the other developers on the team. Update the CI/CD pipeline to run the test scripts.
- D. Create sample events based on the Lambda documentation. Create a Docker container from the Node.js base image to invoke the Lambda functions. Check the response. Document how to run the Docker container for the other developers on the team. Update the CI/CD pipeline to run the Docker container.

Correct Answer: C

Community vote distribution

C (100%)

65703c1 2 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 3 times

KarBiswa 4 months, 2 weeks ago

Selected Answer: C

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/using-sam-cli-local-generate-event.html>

upvoted 3 times

SerialiDr 6 months, 2 weeks ago

Selected Answer: C

This is the most suitable option. AWS SAM CLI is a tool designed for building, testing, and debugging serverless applications, and it includes the ability to locally test Lambda functions. The `sam local invoke` command allows you to invoke Lambda functions locally in a Docker container that simulates the Lambda execution environment. The `sam local generate-event` command can be used to generate sample events. This approach allows developers to run tests locally and can be integrated into CI/CD pipelines.

upvoted 3 times

dilleman 9 months, 3 weeks ago

Selected Answer: C

C should be correct

upvoted 2 times

Digo30sp 10 months ago

Selected Answer: C

The correct answer is (C).

Solution (C) is the best option to meet the developer's requirements. The AWS SAM CLI tool provides an easy way to generate sample events and invoke Lambda functions locally. The solution is also easy to document and integrate into the CI/CD pipeline.

upvoted 4 times

A developer is troubleshooting an application that uses Amazon DynamoDB in the us-west-2 Region. The application is deployed to an Amazon EC2 instance. The application requires read-only permissions to a table that is named Cars. The EC2 instance has an attached IAM role that contains the following IAM policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ReadOnlyAPIActions",
      "Effect": "Allow",
      "Action": [
        "dynamodb:GetItem",
        "dynamodb:BatchGetItem",
        "dynamodb:Scan",
        "dynamodb:Query",
        "dynamodb:ConditionCheckItem"
      ],
      "Resource": "arn:aws:dynamodb:us-west-2:account-id:table/Cars"
    }
  ]
}
```

When the application tries to read from the Cars table, an Access Denied error occurs.

How can the developer resolve this error?


- A. Modify the IAM policy resource to be "arn:aws:dynamodb:us-west-2:account-id:table/*".
- B. Modify the IAM policy to include the dynamodb:* action.
- C. Create a trust policy that specifies the EC2 service principal. Associate the role with the policy.
- D. Create a trust relationship between the role and dynamodb.amazonaws.com.

Correct Answer: C

Community vote distribution

C (85%)

Other

 **LemonGremlin** Highly Voted 1 year, 3 months ago

Selected Answer: C

The most reasonable answer here is C. But I think the question is missing some information.

<https://aws.amazon.com/blogs/security/how-to-use-trust-policies-with-iam-roles/>

upvoted 6 times

 **koniczny69** 12 months ago

What is a trust policy?

I know trust relationship, not a trust policy.

upvoted 1 times

 **albert_kuo** Most Recent 4 months ago

Selected Answer: A

The current policy specifies the resource as "arn:aws:dynamodb:us-west-2:account-id:table/Cars". This means the policy only applies to the specific "Cars" table.



However, when working with DynamoDB, it's common for the SDK or API calls to need permissions on the table's indexes as well as the table itself.

By changing the resource to "arn:aws:dynamodb:us-west-2:account-id:table/*", you're granting the specified read-only actions (GetItem, BatchGetItem, Scan, Query, ConditionCheckItem) on all tables in the account in the us-west-2 region, including any indexes associated with the

Cars table.

This broader permission will likely resolve the Access Denied error, as it will cover both the table and any associated indexes.



upvoted 1 times

  **wh1t4k3r** 5 months, 1 week ago

Selected Answer: C

Well, I will guess that this question is badly written on purpose. Anyway: C makes more sense since A and B are going against best practices of least privilege. D makes no sense since the role must trust the service that will use it rather than the service that will be accessed.

upvoted 1 times

  **65703c1** 8 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

  **PrakashM14** 1 year, 3 months ago

Selected Answer: D

D.Create a trust relationship between the role and dynamodb.amazonaws.com.

Explanation:

Trust Relationship: In AWS, a trust relationship defines who or what entity can assume a role. In this case, the role attached to the EC2 instance needs to trust DynamoDB. The trust relationship is specified in a JSON policy document.

DynamoDB Service Principal: The correct service principal for DynamoDB is dynamodb.amazonaws.com. This is the entity that the role needs to trust to allow access to DynamoDB resources.

upvoted 1 times

  **konieczny69** 12 months ago

Complete nonsense. Role needs to trust EC2, since its the EC2 who is to assume the role.

upvoted 3 times

  **Digo30sp** 1 year, 3 months ago

Selected Answer: C

<https://www.examttopics.com/discussions/amazon/view/96497-exam-aws-certified-developer-associate-topic-1-question-380/>

upvoted 3 times

When using the AWS Encryption SDK, how does the developer keep track of the data encryption keys used to encrypt data?

- A. The developer must manually keep track of the data encryption keys used for each data object.
- B. The SDK encrypts the data encryption key and stores it (encrypted) as part of the returned ciphertext.
- C. The SDK stores the data encryption keys automatically in Amazon S3.
- D. The data encryption key is stored in the Userdata for the EC2 instance.

Correct Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **tsangckl** 7 months, 2 weeks ago

This appear at 17 Jun exam
upvoted 1 times

🗳️ 👤 **gabyslim** 5 months ago

What was the correct answer to it?
upvoted 1 times

🗳️ 👤 **65703c1** 8 months, 1 week ago

Selected Answer: B

B is the correct answer.
upvoted 3 times

🗳️ 👤 **KarBiswa** 11 months ago

Selected Answer: B

<https://docs.aws.amazon.com/encryption-sdk/latest/developer-guide/concepts.html#:~:text=An%20encryption%20context%20is%20a,encrypted%20message%20that%20it%20returns.>
upvoted 2 times

🗳️ 👤 **SerialiDr** 1 year ago

Selected Answer: B

In the AWS Encryption SDK, data is encrypted using a DEK. This DEK is then encrypted with a Key Encryption Key (KEK), usually managed by AWS Key Management Service (AWS KMS) or another key management infrastructure. The encrypted DEK is stored alongside the encrypted data (ciphertext). This allows the SDK to manage the DEKs seamlessly.
upvoted 3 times

🗳️ 👤 **TanTran04** 1 year, 1 month ago

Selected Answer: B

When using the AWS Encryption SDK, it is a common practice to encrypt the data encryption key (DEK) along with the data. The DEK is used to encrypt the actual data, and it is itself encrypted using a key management system, often called a key encryption key (KEK). This encrypted DEK is then stored alongside the encrypted data.
upvoted 3 times

🗳️ 👤 **dilleman** 1 year, 3 months ago

Selected Answer: B

B is correct
upvoted 3 times

🗳️ 👤 **Digo30sp** 1 year, 3 months ago

Selected Answer: B

<https://www.examttopics.com/discussions/amazon/view/96427-exam-aws-certified-developer-associate-topic-1-question-398/>
upvoted 3 times

An application that runs on AWS Lambda requires access to specific highly confidential objects in an Amazon S3 bucket. In accordance with the principle of least privilege, a company grants access to the S3 bucket by using only temporary credentials.

How can a developer configure access to the S3 bucket in the MOST secure way?

- A. Hardcode the credentials that are required to access the S3 objects in the application code. Use the credentials to access the required S3 objects.
- B. Create a secret access key and access key ID with permission to access the S3 bucket. Store the key and key ID in AWS Secrets Manager. Configure the application to retrieve the Secrets Manager secret and use the credentials to access the S3 objects.
- C. Create a Lambda function execution role. Attach a policy to the role that grants access to specific objects in the S3 bucket.
- D. Create a secret access key and access key ID with permission to access the S3 bucket. Store the key and key ID as environment variables in Lambda. Use the environment variables to access the required S3 objects.

Correct Answer: C

Community vote distribution

C (76%)

B (24%)

 **dilleman** Highly Voted 1 year, 3 months ago

Selected Answer: C

C should be correct:

<https://docs.aws.amazon.com/lambda/latest/operatorguide/least-privilege.html>

upvoted 12 times

 **Digo30sp** Highly Voted 1 year, 3 months ago

Selected Answer: B

The correct answer is (B).

Option (B) is the most secure way to configure S3 bucket access because the credentials are stored in a safe and secure location. AWS Secrets Manager uses public key cryptography to protect stored secrets.

upvoted 6 times

 **dezoito** 1 year, 3 months ago

B goes against the least privilege principle because it gives access to the whole bucket

upvoted 6 times

 **aws_god** 3 months ago

It states: "grants access to specific objects in the S3 bucket" so it will not give access to the whole bucket

upvoted 1 times

 **Melisa202401** 9 months, 4 weeks ago

Store credentials in aws secret manager, it will be rotated => so it comply the least privilege principle!

upvoted 1 times

 **65703c1** Most Recent 8 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

 **SerialIDr** 1 year ago

Selected Answer: C

This is the most secure and recommended approach. By attaching an IAM policy to the Lambda execution role that grants access only to the specific S3 objects needed, you adhere to the principle of least privilege. This method also uses AWS's built-in mechanism for providing temporary credentials to the Lambda function, eliminating the need to manage access keys.

upvoted 2 times

 **LemonGremlin** 1 year, 3 months ago

Selected Answer: C

C. Create a Lambda function execution role. Attach a policy to the role that grants access to specific objects in the S3 bucket.
upvoted 4 times

A developer has code that is stored in an Amazon S3 bucket. The code must be deployed as an AWS Lambda function across multiple accounts in the same AWS Region as the S3 bucket. An AWS CloudFormation template that runs for each account will deploy the Lambda function.

What is the MOST secure way to allow CloudFormation to access the Lambda code in the S3 bucket?

- A. Grant the CloudFormation service role the S3 ListBucket and GetObject permissions. Add a bucket policy to Amazon S3 with the principal of "AWS": [account numbers].
- B. Grant the CloudFormation service role the S3 GetObject permission. Add a bucket policy to Amazon S3 with the principal of "*".
- C. Use a service-based link to grant the Lambda function the S3 ListBucket and GetObject permissions by explicitly adding the S3 bucket's account number in the resource.
- D. Use a service-based link to grant the Lambda function the S3 GetObject permission. Add a resource of "*" to allow access to the S3 bucket.

Correct Answer: A

Community vote distribution

A (92%)

8%

🗨️ **NSA_Poker** 4 months, 2 weeks ago

Selected Answer: B

(CD) eliminated. service-based link is not supported by Lambda.

(A) S3 ListBucket permission violates the principle of least privilege and therefore is not the most secure. Bucket policy to list principles of multiple accounts requires additional overhead. The list can change.

(B) allows the CloudFormation service role to access the S3 bucket from any account, as long as it has the S3 GetObject permission. The bucket policy grants access to any principal with the GetObject permission, which is the least privilege needed to deploy the Lambda code.

upvoted 1 times

🗨️ **65703c1** 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ **SerialiDr** 1 year ago

Selected Answer: A

This approach is secure and provides a granular level of control. By granting the CloudFormation service role in each account the necessary S3 permissions and specifying the account numbers in the S3 bucket policy, you ensure that only the specified accounts can access the Lambda code. However, the ListBucket permission is not necessary if the CloudFormation template already knows the exact S3 object key.

upvoted 4 times

🗨️ **TanTran04** 1 year, 1 month ago

Selected Answer: A

Following ChatGPT 3.5, Option A is the best choice. I guess.

- Follows the principle of least privilege by granting only the necessary permissions (ListBucket and GetObject) to the CloudFormation service role.

- Adding a bucket policy with the principal of "AWS": [account numbers] restricts access to only the specified AWS accounts, providing a more secure access control mechanism.

- This ensures that only the CloudFormation service role in the specified AWS accounts can access the Lambda code in the S3 bucket.

upvoted 2 times

🗨️ **Digo30sp** 1 year, 3 months ago

Selected Answer: A

The correct answer is (A).

Option (A) is the safest way to allow CloudFormation to access the Lambda code in the S3 bucket because it limits access to the specific

accounts that need to deploy the Lambda functions. The bucket policy grants S3 ListBucket and GetObject permissions to the CloudFormation service role only for the accounts specified in the principal.

upvoted 4 times

A developer at a company needs to create a small application that makes the same API call once each day at a designated time. The company does not have infrastructure in the AWS Cloud yet, but the company wants to implement this functionality on AWS.

Which solution meets these requirements in the MOST operationally efficient manner?

- A. Use a Kubernetes cron job that runs on Amazon Elastic Kubernetes Service (Amazon EKS).
- B. Use an Amazon Linux crontab scheduled job that runs on Amazon EC2.
- C. Use an AWS Lambda function that is invoked by an Amazon EventBridge scheduled event.
- D. Use an AWS Batch job that is submitted to an AWS Batch job queue.

Correct Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **65703c1** 2 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 2 times

🗨️ 👤 **joshnort** 6 months ago

Selected Answer: C

<https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-create-rule-schedule.html>

upvoted 3 times

🗨️ 👤 **dilleman** 9 months, 3 weeks ago

Selected Answer: C

C is correct

upvoted 3 times

🗨️ 👤 **Digo30sp** 10 months ago

Selected Answer: C

<https://www.examttopics.com/discussions/amazon/view/88703-exam-aws-certified-developer-associate-topic-1-question-229/>

upvoted 2 times

A developer is building a serverless application that is based on AWS Lambda. The developer initializes the AWS software development kit (SDK) outside of the Lambda handler function.

What is the PRIMARY benefit of this action?

- A. Improves legibility and stylistic convention
- B. Takes advantage of runtime environment reuse
- C. Provides better error handling
- D. Creates a new SDK instance for each invocation

Correct Answer: B

Community vote distribution

B (100%)

🗨️ **tsangckl** 7 months, 2 weeks ago

This appear at 17 Jun exam
upvoted 1 times

🗨️ **65703c1** 8 months, 1 week ago

Selected Answer: B

B is the correct answer.
upvoted 2 times

🗨️ **dilleman** 1 year, 3 months ago

Selected Answer: B

B it is!
upvoted 3 times

🗨️ **Digo30sp** 1 year, 3 months ago

Selected Answer: B

The correct answer is (B).

Initializing the AWS SDK outside of the Lambda handler function takes advantage of runtime environment reuse. This means that the SDK only needs to be initialized once for all Lambda function invocations. This can improve application performance and efficiency.

upvoted 3 times

A company is using Amazon RDS as the backend database for its application. After a recent marketing campaign, a surge of read requests to the database increased the latency of data retrieval from the database. The company has decided to implement a caching layer in front of the database. The cached content must be encrypted and must be highly available.

Which solution will meet these requirements?

- A. Amazon CloudFront
- B. Amazon ElastiCache for Memcached
- C. Amazon ElastiCache for Redis in cluster mode
- D. Amazon DynamoDB Accelerator (DAX)

Correct Answer: C

Community vote distribution

C (100%)

  **dilleman** Highly Voted 9 months, 3 weeks ago

Selected Answer: C

Should be C since ElastiCache for Redis supports encryption at rest and in transit. ElastiCache for Memcached does not support encryption at rest.

DynamoDB Accelerator is for DynamoDB and does not fit this case.



upvoted 11 times

  **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

  **SerialiDr** 6 months, 2 weeks ago

Selected Answer: C

ElastiCache for Redis provides both encryption in transit and at rest. In cluster mode, it also offers high availability and scalability. This makes it well-suited for caching database queries while ensuring data security and high availability.

upvoted 3 times

  **Digo30sp** 10 months ago

Selected Answer: C

<https://www.examttopics.com/discussions/amazon/view/82917-exam-aws-certified-developer-associate-topic-1-question-95/>

upvoted 2 times

A developer at a company recently created a serverless application to process and show data from business reports. The application's user interface (UI) allows users to select and start processing the files. The UI displays a message when the result is available to view. The application uses AWS Step Functions with AWS Lambda functions to process the files. The developer used Amazon API Gateway and Lambda functions to create an API to support the UI.

The company's UI team reports that the request to process a file is often returning timeout errors because of the size or complexity of the files. The UI team wants the API to provide an immediate response so that the UI can display a message while the files are being processed. The backend process that is invoked by the API needs to send an email message when the report processing is complete.

What should the developer do to configure the API to meet these requirements?

- A. Change the API Gateway route to add an X-Amz-Invocation-Type header with a static value of 'Event' in the integration request. Deploy the API Gateway stage to apply the changes.
- B. Change the configuration of the Lambda function that implements the request to process a file. Configure the maximum age of the event so that the Lambda function will run asynchronously.
- C. Change the API Gateway timeout value to match the Lambda function timeout value. Deploy the API Gateway stage to apply the changes.
- D. Change the API Gateway route to add an X-Amz-Target header with a static value of 'Async' in the integration request. Deploy the API Gateway stage to apply the changes.

Correct Answer: A

Community vote distribution

A (75%)

D (25%)

 **SerialiDr** Highly Voted 1 year ago

Selected Answer: A

Change the API Gateway route to add an X-Amz-Invocation-Type header with a static value of 'Event' in the integration request: This is the correct approach. By setting the X-Amz-Invocation-Type header to Event in the API Gateway integration request, the API Gateway will invoke the Lambda function asynchronously. In asynchronous execution, the Lambda function returns an immediate response (202 or Accepted status) to API Gateway, which can then relay it back to the UI. Meanwhile, the Lambda function processes the file in the background.

upvoted 5 times

 **albert_kuo** Most Recent 4 months ago

Selected Answer: A

```
curl -X POST \
-H "X-Amz-Invocation-Type: Event" \
https://lambda-endpoint.amazonaws.com/2015-03-31/functions/my-function/invocations \
-d '{"key1": "value1"}
```


upvoted 1 times

 **Itfalcon** 4 months, 2 weeks ago

According chatgpt correct is A. D is wrong because header value is set for Async and it should be set for Event

This header is incorrect for this use case. The correct header for asynchronous invocation of Lambda is X-Amz-Invocation-Type with the value 'Event', not 'Async'.


upvoted 1 times

 **Saurabh04** 5 months, 2 weeks ago

Selected Answer: D

By adding the X-Amz-Target header with a value of 'Async', the API Gateway will invoke the backend Lambda function asynchronously.

upvoted 1 times

 **65703c1** 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ **JLLNOR** 1 year, 1 month ago

Selected Answer: A

<https://docs.aws.amazon.com/apigateway/latest/developerguide/set-up-lambda-integration-async.html>

upvoted 4 times

🗨️ **Certified101** 1 year, 1 month ago

Selected Answer: A

<https://www.examttopics.com/discussions/amazon/view/82655-exam-aws-certified-developer-associate-topic-1-question-85/>

upvoted 3 times

🗨️ **TanTran04** 1 year, 1 month ago

Selected Answer: D

Option A is incorrect because the X-Amz-Invocation-Type header with a static value of 'Event' is used for the AWS Lambda asynchronous invocation, but it doesn't address the issue of providing an immediate response to the UI.

Option D is the correct choice. By adding an X-Amz-Target header with a static value of 'Async' in the integration request, the API Gateway will immediately return a response to the UI, allowing it to display a message while the backend processing continues asynchronously. This ensures that the UI team does not encounter timeout errors due to long-running processes.

upvoted 1 times

🗨️ **TanTran04** 1 year, 1 month ago

I miss something, Option D is undocumented.

=> A is the best choice

upvoted 2 times

🗨️ **LemonGremlin** 1 year, 3 months ago

Selected Answer: A

Reference: <https://docs.aws.amazon.com/apigateway/latest/developerguide/set-up-lambda-integration-async.html>

upvoted 2 times

🗨️ **kashtelyan** 1 year, 3 months ago

Selected Answer: D

Option A involves changing the API Gateway route to add an X-Amz-Invocation-Type header with a static value of 'Event' in the integration request. This header is typically used when you want to invoke a Lambda function asynchronously, but it doesn't ensure that you get an immediate response. It essentially sends the request to a queue for asynchronous execution and doesn't wait for the processing to complete before providing a response.

In contrast, option D suggests using the X-Amz-Target header with a static value of 'Async,' which is a more appropriate choice when you need to provide an immediate response to the client while offloading the processing for background execution. This approach better aligns with the requirement of displaying a message to the user while the files are being processed, which is typically achieved through asynchronous processing with notification upon completion.

upvoted 4 times

🗨️ **Digo30sp** 1 year, 3 months ago

Selected Answer: A

A) <https://www.examttopics.com/discussions/amazon/view/82655-exam-aws-certified-developer-associate-topic-1-question-85/>

upvoted 2 times

🗨️ **fordiscussiontwo** 1 year, 3 months ago

aaaaaaaaAAAAAAAAAAAAAAAAAAAA

upvoted 2 times

A developer has an application that is composed of many different AWS Lambda functions. The Lambda functions all use some of the same dependencies. To avoid security issues, the developer is constantly updating the dependencies of all of the Lambda functions. The result is duplicated effort for each function.

How can the developer keep the dependencies of the Lambda functions up to date with the LEAST additional complexity?

- A. Define a maintenance window for the Lambda functions to ensure that the functions get updated copies of the dependencies.
- B. Upgrade the Lambda functions to the most recent runtime version.
- C. Define a Lambda layer that contains all of the shared dependencies.
- D. Use an AWS CodeCommit repository to host the dependencies in a centralized location.

Correct Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **65703c1** 2 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗨️ 👤 **TanTran04** 7 months, 2 weeks ago

Selected Answer: C

To share dependencies across multiple functions. After you create a layer, you can apply it to any number of functions in your account. Without layers, you need to include the same dependencies in each individual deployment package.

<https://docs.aws.amazon.com/lambda/latest/dg/chapter-layers.html>

upvoted 2 times

🗨️ 👤 **dilleman** 9 months, 3 weeks ago

Selected Answer: C

C is correct.

upvoted 2 times

🗨️ 👤 **Digo30sp** 10 months ago

Selected Answer: C

C) <https://www.examttopics.com/discussions/amazon/view/96245-exam-aws-certified-developer-associate-topic-1-question-436/>

upvoted 2 times

A mobile app stores blog posts in an Amazon DynamoDB table. Millions of posts are added every day, and each post represents a single item in the table. The mobile app requires only recent posts. Any post that is older than 48 hours can be removed.

What is the MOST cost-effective way to delete posts that are older than 48 hours?

- A. For each item, add a new attribute of type String that has a timestamp that is set to the blog post creation time. Create a script to find old posts with a table scan and remove posts that are older than 48 hours by using the BatchWriteItem API operation. Schedule a cron job on an Amazon EC2 instance once an hour to start the script.
- B. For each item, add a new attribute of type String that has a timestamp that is set to the blog post creation time. Create a script to find old posts with a table scan and remove posts that are older than 48 hours by using the BatchWriteItem API operation. Place the script in a container image. Schedule an Amazon Elastic Container Service (Amazon ECS) task on AWS Fargate that invokes the container every 5 minutes.
- C. For each item, add a new attribute of type Date that has a timestamp that is set to 48 hours after the blog post creation time. Create a global secondary index (GSI) that uses the new attribute as a sort key. Create an AWS Lambda function that references the GSI and removes expired items by using the BatchWriteItem API operation. Schedule the function with an Amazon CloudWatch event every minute.
- D. For each item, add a new attribute of type Number that has a timestamp that is set to 48 hours after the blog post creation time. Configure the DynamoDB table with a TTL that references the new attribute.

Correct Answer: D

Community vote distribution

D (100%)

65703c1 2 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 2 times

SerialIDr 6 months, 2 weeks ago

Selected Answer: D

This is the most cost-effective and efficient solution. The TTL feature allows DynamoDB to automatically delete items past a certain timestamp, which is perfect for this use case. By adding a TTL attribute to each item (set to 48 hours after the post creation time), DynamoDB will automatically delete the items when they expire, without any need for custom scripts, additional AWS services, or manual intervention.

upvoted 2 times

dilleman 9 months, 3 weeks ago

Selected Answer: D

D is correct. DynamoDB tables can clean up data itself based on provided configuration.

upvoted 3 times

Digo30sp 10 months ago

Selected Answer: D

The correct answer is (D).

Solution (D) is the most cost-effective because it uses DynamoDB's Time to Live (TTL) to automatically remove expired items. The TTL is an item attribute that specifies the duration of time that an item should remain in the table. When an item's TTL expires, the item is automatically deleted from the table.

upvoted 3 times

A developer is modifying an existing AWS Lambda function. While checking the code, the developer notices hardcoded parameter values for an Amazon RDS for SQL Server user name, password, database, host, and port. There are also hardcoded parameter values for an Amazon DynamoDB table, an Amazon S3 bucket, and an Amazon Simple Notification Service (Amazon SNS) topic.

The developer wants to securely store the parameter values outside the code in an encrypted format and wants to turn on rotation for the credentials. The developer also wants to be able to reuse the parameter values from other applications and to update the parameter values without modifying code.

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

- A. Create an RDS database secret in AWS Secrets Manager. Set the user name, password, database, host, and port. Turn on secret rotation. Create encrypted Lambda environment variables for the DynamoDB table, S3 bucket, and SNS topic.
- B. Create an RDS database secret in AWS Secrets Manager. Set the user name, password, database, host, and port. Turn on secret rotation. Create SecureString parameters in AWS Systems Manager Parameter Store for the DynamoDB table, S3 bucket, and SNS topic.
- C. Create RDS database parameters in AWS Systems Manager Parameter Store for the user name, password, database, host, and port. Create encrypted Lambda environment variables for the DynamoDB table, S3 bucket, and SNS topic. Create a Lambda function and set the logic for the credentials rotation task. Schedule the credentials rotation task in Amazon EventBridge.
- D. Create RDS database parameters in AWS Systems Manager Parameter Store for the user name, password, database, host, and port. Store the DynamoDB table, S3 bucket, and SNS topic in Amazon S3. Create a Lambda function and set the logic for the credentials rotation. Invoke the Lambda function on a schedule.

Correct Answer: B

Community vote distribution

B (100%)

65703c1 2 months, 1 week ago

Selected Answer: B

B is the correct answer.
upvoted 1 times

SerialiDr 6 months, 2 weeks ago

Selected Answer: B

AWS Secrets Manager is designed to handle sensitive information like database credentials and supports automatic rotation. Using SecureString parameters in Systems Manager Parameter Store for other parameters provides a secure and centralized way to manage them. This approach also enables reusability and easy updating without code modifications.
upvoted 4 times

dilleman 9 months, 3 weeks ago

Selected Answer: B

B is correct
upvoted 3 times

Digo30sp 10 months ago

Selected Answer: B

B) <https://www.examtactics.com/discussions/amazon/view/88929-exam-aws-certified-developer-associate-topic-1-question-338/>
upvoted 3 times

A developer accesses AWS CodeCommit over SSH. The SSH keys configured to access AWS CodeCommit are tied to a user with the following permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "codecommit:BatchGetRepositories",
        "codecommit:Get*",
        "codecommit:List*",
        "codecommit:GitPull"
      ],
      "Resource": "*"
    }
  ]
}
```

The developer needs to create/delete branches.

Which specific IAM permissions need to be added, based on the principle of least privilege?

- A. "codecommit:CreateBranch"
"codecommit>DeleteBranch"
- B. "codecommit:Put*"
- C. "codecommit:Update*"
- D. "codecommit:*"

Correct Answer: A

Community vote distribution

A (100%)

65703c1 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 3 times

TanTran04 7 months, 2 weeks ago

Selected Answer: A

Take a look at CodeCommit API Operations and Required Permissions for Actions on Branches in

<https://docs.aws.amazon.com/codecommit/latest/userguide/auth-and-access-control-permissions-reference.html#aa-branches> => A

upvoted 2 times

dilleman 9 months, 3 weeks ago

Selected Answer: A

A of course

upvoted 3 times

Digo30sp 10 months ago

Selected Answer: A

A) <https://www.examtactics.com/discussions/amazon/view/4364-exam-aws-certified-developer-associate-topic-1-question-190/>

upvoted 2 times

An application that is deployed to Amazon EC2 is using Amazon DynamoDB. The application calls the DynamoDB REST API. Periodically, the application receives a ProvisionedThroughputExceededException error when the application writes to a DynamoDB table.

Which solutions will mitigate this error MOST cost-effectively? (Choose two.)

- A. Modify the application code to perform exponential backoff when the error is received.
- B. Modify the application to use the AWS SDKs for DynamoDB.
- C. Increase the read and write throughput of the DynamoDB table.
- D. Create a DynamoDB Accelerator (DAX) cluster for the DynamoDB table.
- E. Create a second DynamoDB table. Distribute the reads and writes between the two tables.

Correct Answer: AB

Community vote distribution

AB (94%) 6%

 **SerialiDr** Highly Voted 1 year ago

Selected Answer: AB

A. Modify the application code to perform exponential backoff when the error is received: This is a cost-effective and recommended approach. Exponential backoff is a standard error-retry strategy where the time between retries gradually increases. This strategy helps to efficiently manage request retries without immediately consuming additional throughput, thus reducing the likelihood of repeatedly hitting the throughput limits.

B. Modify the application to use the AWS SDKs for DynamoDB: The AWS SDKs implement best practices, including automatic retry logic with exponential backoff. Using an AWS SDK for DynamoDB can simplify the implementation and is more efficient than directly calling the DynamoDB REST API. This change can help mitigate throughput exceedance errors.

upvoted 6 times

 **f271c23** Most Recent 1 month, 3 weeks ago

Selected Answer: AB

Yes AB is the right set of answers.

upvoted 1 times

 **CloudChingon** 2 months, 2 weeks ago

Key Words Why They Matter

ProvisionedThroughputExceededException Indicates the application is exceeding the provisioned capacity of the DynamoDB table.

Periodically Suggests that the issue happens under specific conditions, such as traffic spikes.

REST API Implies the application is making direct API calls, possibly without SDK optimizations.

Cost-effectively Points towards solutions that avoid increasing DynamoDB capacity (which costs more).

Writes to a DynamoDB table

Identifies the operation likely causing the throttling issue.

Modify application code Indicates that a code change is necessary to handle throttling errors better.

Exponential backoff A retry strategy to reduce the impact of throttling.

AWS SDKs Built-in solutions for efficient, optimized interactions with DynamoDB.

upvoted 1 times

 **aws_god** 3 months ago



Selected Answer: BC

This is what the exception looks like:

```
"boto.dynamodb2.exceptions.ProvisionedThroughputExceededException: ProvisionedThroughputExceededException: 400 Bad Request
{'u:message': u'The level of configured provisioned throughput for the table was exceeded. Consider increasing your provisioning level with the
UpdateTable API', u'_type': u'com.amazonaws.dynamodb.v20120810#ProvisionedThroughputExceededException}'"
```

You need B since the AWS SDK includes automatic retry logic with exponential backoff specifically designed to deal with this automatically and C since increasing the provisioned level also helps solve this issue.

upvoted 1 times

  **65703c1** 8 months, 1 week ago

Selected Answer: AB

AB is the correct answer.

upvoted 1 times

  **TanTran04** 1 year, 1 month ago

Selected Answer: AC

Following ChatGPT 3.5, Option A and C

Option B (Modify the application to use the AWS SDKs for DynamoDB) is not directly related to resolving throughput issues. It's generally recommended to use the AWS SDKs as they provide more efficient and convenient ways to interact with AWS services, but it may not directly address the ProvisionedThroughputExceededException issue.

upvoted 1 times

  **dilleman** 1 year, 3 months ago

Selected Answer: AB

A and B.

Exponential backoff is a standard error-handling strategy for network applications. The idea is to retry a failed request with increasing delays between each attempt.

And the AWS SDKs have built-in support for handling these errors.

upvoted 4 times

  **Digo30sp** 1 year, 3 months ago

Selected Answer: AB

A and B: <https://www.examttopics.com/discussions/amazon/view/69199-exam-aws-certified-developer-associate-topic-1-question-385/>

upvoted 4 times

When a developer tries to run an AWS CodeBuild project, it raises an error because the length of all environment variables exceeds the limit for the combined maximum of characters.

What is the recommended solution?

- A. Add the `export LC_ALL="en_US.utf8"` command to the `pre_build` section to ensure POSIX localization.
- B. Use Amazon Cognito to store key-value pairs for large numbers of environment variables.
- C. Update the settings for the build project to use an Amazon S3 bucket for large numbers of environment variables.
- D. Use AWS Systems Manager Parameter Store to store large numbers of environment variables.

Correct Answer: D

Community vote distribution

D (100%)

🗨️ **65703c1** 2 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

🗨️ **SerialiDr** 6 months, 2 weeks ago

Selected Answer: D

AWS Systems Manager Parameter Store is specifically designed for managing configuration data and secrets. It can store large numbers of parameters, including environment variables, and makes them easily accessible and manageable. It also provides features like versioning, fine-grained access control, and integration with AWS Identity and Access Management (IAM).

upvoted 4 times

🗨️ **dilleman** 9 months, 3 weeks ago

Selected Answer: D

Best solution is D

upvoted 4 times

🗨️ **Digo30sp** 10 months ago

Selected Answer: D

D) <https://docs.aws.amazon.com/codebuild/latest/userguide/troubleshooting.html>

upvoted 4 times

A company is expanding the compatibility of its photo-sharing mobile app to hundreds of additional devices with unique screen dimensions and resolutions. Photos are stored in Amazon S3 in their original format and resolution. The company uses an Amazon CloudFront distribution to serve the photos. The app includes the dimension and resolution of the display as GET parameters with every request.

A developer needs to implement a solution that optimizes the photos that are served to each device to reduce load time and increase photo quality.

Which solution will meet these requirements MOST cost-effectively?

- A. Use S3 Batch Operations to invoke an AWS Lambda function to create new variants of the photos with the required dimensions and resolutions. Create a dynamic CloudFront origin that automatically maps the request of each device to the corresponding photo variant.
- B. Use S3 Batch Operations to invoke an AWS Lambda function to create new variants of the photos with the required dimensions and resolutions. Create a Lambda@Edge function to route requests to the corresponding photo variant by using request headers.
- C. Create a Lambda@Edge function that optimizes the photos upon request and returns the photos as a response. Change the CloudFront TTL cache policy to the maximum value possible.
- D. Create a Lambda@Edge function that optimizes the photos upon request and returns the photos as a response. In the same function, store a copy of the processed photos on Amazon S3 for subsequent requests.

Correct Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **tsangckl** 7 months, 2 weeks ago

This appear at 17 Jun exam
upvoted 1 times

🗳️ 👤 **65703c1** 8 months, 1 week ago

Selected Answer: D

D is the correct answer.
upvoted 2 times

🗳️ 👤 **Melisa202401** 10 months ago

C
I dont see any aws docs about Cloufront cache=> so maybe it is cost-effective
upvoted 1 times

🗳️ 👤 **joshnort** 12 months ago

Selected Answer: D

<https://aws.amazon.com/blogs/networking-and-content-delivery/resizing-images-with-amazon-cloudfront-lambdaedge-aws-cdn-blog/>
upvoted 3 times

🗳️ 👤 **SerialiDr** 1 year ago

Selected Answer: D

This solution is the most cost-effective. Lambda@Edge processes the photos dynamically based on the device's requirements, which means no pre-generation of multiple variants is required. Processed photos are stored on S3, ensuring that subsequent requests for the same photo variant are served directly from S3, reducing Lambda@Edge invocations and further optimizing costs.
upvoted 4 times

🗳️ 👤 **Mimi666** 1 year, 1 month ago

Selected Answer: D

<https://aws.amazon.com/es/blogs/networking-and-content-delivery/image-optimization-using-amazon-cloudfront-and-aws-lambda/>
upvoted 4 times

🗳️ 👤 **jingle4944** 1 year, 2 months ago

According to <https://aws.amazon.com/blogs/networking-and-content-delivery/resizing-images-with-amazon-cloudfront-lambdaedge-aws-cdn-blog/>, "static resources like images should have a long Time to Live (TTL) as possible to improve cache-hit ratios.". The photo cache here is likely

to be static and should be preserved forever.

upvoted 1 times

🗨️ 👤 **BrainFried** 6 months, 1 week ago

CloudFront has a Maximum TTL of 365 days. Would it not be cheaper to store the images in the CloudFront cache, instead of storing it in S3 which would incur costs?

We may need to assume it would be unlikely the users would access the same photo more than a year after the initial access.

upvoted 1 times

🗨️ 👤 **ut18** 1 year, 3 months ago

Why not B?

The developer can use S3 Batch Operations to create new variants of the photos with the required dimensions and resolutions.

upvoted 1 times

🗨️ 👤 **TallManDan** 1 year, 3 months ago

Selected Answer: D

You only want to convert the pictures that get requests. If you convert them all through batch processing, you have wasted time and expense on any possible photo that never gets viewed. The Minimum TTL is set to 60 seconds, the Default TTL is set to 300 seconds, and the Maximum TTL is set to 3600 seconds. S3 is the way to go.

upvoted 2 times

🗨️ 👤 **BrainFried** 6 months, 1 week ago

CloudFront cache has a Maximum TTL of 365 days. Would it not be cheaper to store the images in the CloudFront cache, instead of storing it in S3 which would incur costs?

We may need to assume it would be unlikely the users would access the same photo more than a year after the initial access.

upvoted 1 times

🗨️ 👤 **dilleman** 1 year, 3 months ago

Selected Answer: D

D is correct

upvoted 1 times

🗨️ 👤 **Digo30sp** 1 year, 3 months ago

Selected Answer: D

D) <https://www.examttopics.com/discussions/amazon/view/89564-exam-aws-certified-developer-associate-topic-1-question-320/>

upvoted 1 times

A company is building an application for stock trading. The application needs sub-millisecond latency for processing trade requests. The company uses Amazon DynamoDB to store all the trading data that is used to process each trading request.

A development team performs load testing on the application and finds that the data retrieval time is higher than expected. The development team needs a solution that reduces the data retrieval time with the least possible effort.

Which solution meets these requirements?

- A. Add local secondary indexes (LSIs) for the trading data.
- B. Store the trading data in Amazon S3, and use S3 Transfer Acceleration.
- C. Add retries with exponential backoff for DynamoDB queries.
- D. Use DynamoDB Accelerator (DAX) to cache the trading data.

Correct Answer: D

Community vote distribution

D (100%)

65703c1 2 months, 1 week ago

Selected Answer: D

D is the correct answer.
upvoted 2 times

joshnort 6 months ago

Selected Answer: D

<https://aws.amazon.com/dynamodb/dax/>
upvoted 1 times

SerialiDr 6 months, 2 weeks ago

Selected Answer: D

DAX is an in-memory cache for DynamoDB that delivers fast read performance for your tables at scale by enabling you to get sub-millisecond response times for accessing your data. DAX is particularly beneficial for read-heavy and bursty workloads. Since it reduces the time to retrieve data, it's the most appropriate solution for achieving sub-millisecond latency in data retrieval.
upvoted 4 times

JohnPI 6 months, 2 weeks ago

Selected Answer: D

Use DynamoDB Accelerator (DAX)
upvoted 2 times

TanTran04 7 months, 2 weeks ago

Selected Answer: D

<https://aws.amazon.com/dynamodb/dax/>
DAX delivers up to a 10 times performance improvement—from milliseconds to microseconds—even at millions of requests per second. only pay for the capacity you provision.
upvoted 2 times

dilleman 9 months, 3 weeks ago

Selected Answer: D

This is a perfect scenario for DAX so correct answer is D
upvoted 3 times

Digo30sp 10 months ago

Selected Answer: D

D) <https://www.examttopics.com/discussions/amazon/view/4971-exam-aws-certified-developer-associate-topic-1-question-14/>
upvoted 3 times

A developer is working on a Python application that runs on Amazon EC2 instances. The developer wants to enable tracing of application requests to debug performance issues in the code.


Which combination of actions should the developer take to achieve this goal? (Choose two.)

- A. Install the Amazon CloudWatch agent on the EC2 instances.
- B. Install the AWS X-Ray daemon on the EC2 instances.
- C. Configure the application to write JSON-formatted logs to `/var/log/cloudwatch`.
- D. Configure the application to write trace data to `/var/log/xray`.
- E. Install and configure the AWS X-Ray SDK for Python in the application.

Correct Answer: BE

Community vote distribution

BE (100%)

 **Digo30sp** Highly Voted 10 months ago

Selected Answer: BE

The correct answers are (E) and (B).

(E) is the most important action to enable application request tracking using AWS X-Ray. The AWS X-Ray SDK for Python provides a set of APIs that a developer can use to instrument their application code for tracing.


(B) is the second most important action. The AWS X-Ray daemon runs on each EC2 instance and collects application trace data
upvoted 9 times

 **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: BE

BE is the correct answer.

upvoted 1 times

 **SerialiDr** 6 months, 2 weeks ago

Selected Answer: BE

B. Install the AWS X-Ray daemon on the EC2 instances: This is a required step for enabling AWS X-Ray tracing. The X-Ray daemon listens for traffic on UDP port 2000, gathers raw segment data, and relays it to the AWS X-Ray API. This is necessary for collecting and sending trace data from the application to X-Ray.

E. Install and configure the AWS X-Ray SDK for Python in the application: This is a critical step for enabling X-Ray tracing in your Python application. The X-Ray SDK for Python provides classes and methods to collect data about the requests that your application serves, and sends this data to the X-Ray daemon.

upvoted 3 times

 **NinjaCloud** 9 months ago

Answer: E,B

upvoted 3 times

 **dilleman** 9 months, 3 weeks ago

Selected Answer: BE

B and E

upvoted 4 times

A company has an application that runs as a series of AWS Lambda functions. Each Lambda function receives data from an Amazon Simple Notification Service (Amazon SNS) topic and writes the data to an Amazon Aurora DB instance.

To comply with an information security policy, the company must ensure that the Lambda functions all use a single securely encrypted database connection string to access Aurora.

Which solution will meet these requirements?

- A. Use IAM database authentication for Aurora to enable secure database connections for all the Lambda functions.
- B. Store the credentials and read the credentials from an encrypted Amazon RDS DB instance.
- C. Store the credentials in AWS Systems Manager Parameter Store as a secure string parameter.
- D. Use Lambda environment variables with a shared AWS Key Management Service (AWS KMS) key for encryption.

Correct Answer: C

Community vote distribution

C (63%) A (26%) 11%


 **dilleman** Highly Voted 1 year, 3 months ago

Selected Answer: C

C.

AWS Systems Manager Parameter Store offers a more centralized way to manage encrypted secrets across multiple services than Lambda environment variables, making it a better fit for this scenario.

upvoted 12 times

 **ShawnWon** Highly Voted 1 year, 2 months ago

C.

Option A (IAM database authentication) may provide secure authentication, but it doesn't directly address the storage and retrieval of the connection string.

Option B (storing credentials in an encrypted RDS DB instance) might introduce unnecessary complexity and potential security risks.

Option D (Lambda environment variables with a shared AWS KMS key) is a viable option, but using Parameter Store is generally considered a more centralized and managed approach for storing and retrieving sensitive data in AWS.

Therefore, option C is the most appropriate choice for securely managing the database connection string in this scenario.


upvoted 7 times

 **wh1t4k3r** Most Recent 5 months, 1 week ago

Selected Answer: C


C is the best choice. For those that chose A: simply enabling the DB IAM auth does not address the need to use a single secure string. It would require more steps to make this work regarding lambda execution role, iam policies and etc.

upvoted 1 times

 **tsangckl** 7 months, 2 weeks ago

This appear at 17 Jun exam

upvoted 1 times

 **65703c1** 8 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

 **KarBiswa** 11 months ago

Selected Answer: A

<https://aws.amazon.com/blogs/database/iam-role-based-authentication-to-amazon-aurora-from-serverless-applications/>

upvoted 1 times

🗨️ 👤 **KillThemWithKindness** 11 months, 1 week ago

Selected Answer: A

The developer can create an IAM role with permission to connect to Aurora DB instance and attach it to each Lambda function. The developer can also configure Aurora DB instance to use IAM database authentication and enable encryption in transit using SSL certificates.

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.IAMDBAuth.html>

upvoted 2 times

🗨️ 👤 **rimaSamir** 1 year ago

The answer is A.

<https://aws.amazon.com/ru/blogs/database/iam-role-based-authentication-to-amazon-aurora-from-serverless-applications/>

upvoted 2 times

🗨️ 👤 **rimaSamir** 1 year ago

In Amazon Aurora, you can associate the database users with the IAM user and roles.

upvoted 1 times

🗨️ 👤 **SerialiDr** 1 year ago

Selected Answer: C

C. Store the credentials in AWS Systems Manager Parameter Store as a secure string parameter: This is a strong option. Systems Manager Parameter Store provides secure, hierarchical storage for configuration data and secrets. It can store data such as passwords and database connection strings securely, and it integrates with AWS Key Management Service (AWS KMS) for encryption. Lambda functions can then retrieve the connection string securely at runtime.

D. Use Lambda environment variables with a shared AWS Key Management Service (AWS KMS) key for encryption: While Lambda environment variables can be encrypted with AWS KMS and used to store sensitive information like database connection strings, they are not as centrally manageable as Parameter Store. Each Lambda function's environment variables would need to be updated individually if the connection string changes, which is less efficient and more prone to error.

upvoted 3 times

🗨️ 👤 **TallManDan** 1 year, 3 months ago

Selected Answer: A

<https://aws.amazon.com/blogs/database/iam-role-based-authentication-to-amazon-aurora-from-serverless-applications/>

upvoted 4 times

🗨️ 👤 **Digo30sp** 1 year, 3 months ago

Selected Answer: D

The correct answer is (D).

Solution (D) is the best option because it uses Lambda environment variables with an AWS Key Management Service (AWS KMS) shared key for encryption.

upvoted 3 times

A developer is troubleshooting an Amazon API Gateway API. Clients are receiving HTTP 400 response errors when the clients try to access an endpoint of the API.

How can the developer determine the cause of these errors?

- A. Create an Amazon Kinesis Data Firehose delivery stream to receive API call logs from API Gateway. Configure Amazon CloudWatch Logs as the delivery stream's destination.
- B. Turn on AWS CloudTrail Insights and create a trail. Specify the Amazon Resource Name (ARN) of the trail for the stage of the API.
- C. Turn on AWS X-Ray for the API stage. Create an Amazon CloudWatch Logs log group. Specify the Amazon Resource Name (ARN) of the log group for the API stage.
- D. Turn on execution logging and access logging in Amazon CloudWatch Logs for the API stage. Create a CloudWatch Logs log group. Specify the Amazon Resource Name (ARN) of the log group for the API stage.


Correct Answer: D

Community vote distribution

D (76%)

B (18%)

6%

 **walala97** Highly Voted 7 months ago

Selected Answer: D

why not C,X-Ray is more focused on the internal workings and performance of the API rather than the validity or structure of incoming requests. according to the error 400, it should be the client side error like incorrect request syntax, invalid request message framing, or deceptive request routing


upvoted 6 times

 **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

 **yingying920928** 4 months, 2 weeks ago

Selected Answer: D

Changing answer to D:

CloudTrail records API-level events, but it may not capture the payloads, headers, or other details of the requests and responses that are essential for understanding the cause of HTTP 400 response errors.

upvoted 2 times

 **yingying920928** 4 months, 2 weeks ago

Selected Answer: B

Logging events of API -> AWS CloudTrail


upvoted 1 times

 **KarBiswa** 4 months, 4 weeks ago

Selected Answer: B

Cloud Trail logs

upvoted 2 times

 **KarBiswa** 4 months, 2 weeks ago

Switching my vote to D

upvoted 1 times

 **KarBiswa** 5 months ago

Selected Answer: C

Always the rule of thumb is Cloud trail to trace the logging events of API's.

<https://docs.aws.amazon.com/awsccloudtrail/latest/userguide/logging-insights-events-with-cloudtrail.html#:~:text=AWS%20CloudTrail%20Insights,write%20management%20APIs.>

upvoted 1 times

🗨️ 👤 **KarBiswa** 4 months, 4 weeks ago

Sorry its not C its B

upvoted 1 times

🗨️ 👤 **dezoito** 9 months, 2 weeks ago

D according to <https://docs.aws.amazon.com/apigateway/latest/developerguide/set-up-logging.html>

upvoted 3 times

🗨️ 👤 **dilleman** 9 months, 3 weeks ago

Selected Answer: D

D should be correct

upvoted 2 times

🗨️ 👤 **Digo30sp** 10 months ago

Selected Answer: D

D) <https://www.examttopics.com/discussions/amazon/view/88807-exam-aws-certified-developer-associate-topic-1-question-264/>

upvoted 2 times


A company developed an API application on AWS by using Amazon CloudFront, Amazon API Gateway, and AWS Lambda. The API has a minimum of four requests every second. A developer notices that many API users run the same query by using the POST method. The developer wants to cache the POST request to optimize the API resources.

Which solution will meet these requirements?

- A. Configure the CloudFront cache. Update the application to return cached content based upon the default request headers.
- B. Override the cache method in the selected stage of API Gateway. Select the POST method.
- C. Save the latest request response in Lambda /tmp directory. Update the Lambda function to check the /tmp directory.
- D. Save the latest request in AWS Systems Manager Parameter Store. Modify the Lambda function to take the latest request response from Parameter Store.

Correct Answer: B

Community vote distribution



kr5031 Highly Voted 1 year, 3 months ago

Selected Answer: B

A is incorrect, because of

CloudFront always caches responses to GET and HEAD requests. You can also configure CloudFront to cache responses to OPTIONS requests. CloudFront does not cache responses to requests that use the other methods.

(<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/RequestAndResponseBehaviorCustomOrigin.html>)

upvoted 7 times

dilleman 1 year, 3 months ago

I agree, I think B is correct as well looking into it more.

upvoted 2 times

Jing2023 Highly Voted 1 year, 3 months ago

Selected Answer: B

Why A is not correct

Amazon CloudFront does not cache the responses to POST, PUT, DELETE, and PATCH requests – these requests are proxied back to the origin server. You may enable caching for the responses to OPTIONS requests.

upvoted 5 times

Anandesh Most Recent 6 months ago

Selected Answer: B

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-caching.html#override-api-gateway-stage-cache-for-method-cache>

upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

KarBiswa 11 months ago

Selected Answer: B

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-caching.html#:~:text=When%20you%20enable,caching%20is%20disabled.>

upvoted 2 times

SerialiDr 1 year ago

Selected Answer: B

API Gateway allows caching of responses, and you can enable caching for specific methods, including POST. This option is a viable solution as it leverages the built-in capabilities of API Gateway to cache responses. By configuring caching at the API Gateway stage level, the developer can cache responses to POST requests based on defined parameters.

upvoted 5 times

  **dilleman** 1 year, 3 months ago

Selected Answer: A

A is the correct answer here. CloudFront can be configured to cache based on request headers, query strings, and POST request bodies. Option B might work but it does not work by default and it's not an effective way to solve this.

upvoted 1 times

  **Digo30sp** 1 year, 3 months ago

Selected Answer: B

The correct answer is (B).

Solution (B) is the best option because it uses the Amazon API Gateway cache to cache POST requests.

upvoted 2 times

A company is building a microservices application that consists of many AWS Lambda functions. The development team wants to use AWS Serverless Application Model (AWS SAM) templates to automatically test the Lambda functions. The development team plans to test a small percentage of traffic that is directed to new updates before the team commits to a full deployment of the application.

Which combination of steps will meet these requirements in the MOST operationally efficient way? (Choose two.)

- A. Use AWS SAM CLI commands in AWS CodeDeploy to invoke the Lambda functions to test the deployment.
- B. Declare the EventInvokeConfig on the Lambda functions in the AWS SAM templates with OnSuccess and OnFailure configurations.
- C. Enable gradual deployments through AWS SAM templates.
- D. Set the deployment preference type to Canary10Percent30Minutes. Use hooks to test the deployment.
- E. Set the deployment preference type to Linear10PercentEvery10Minutes. Use hooks to test the deployment.

Correct Answer: CD

Community vote distribution

CD (75%)

AD (15%)

5%

 **dilleman** Highly Voted 1 year, 3 months ago

Selected Answer: CD

C and D should be correct.

Given that "The development team plans to test a small percentage of traffic that is directed to new updates before the team commits to a full deployment of the application." then Option D makes more sense than Option E.

upvoted 6 times

 **Anandesh** Most Recent 6 months, 2 weeks ago

Selected Answer: CD

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverless-apps.html>

You need to enable the gradual deployment first then only codeDeploy will invoke lambda function


upvoted 1 times

 **65703c1** 8 months, 1 week ago

Selected Answer: CD

CD is the correct answer.

upvoted 1 times

 **KarBiswa** 10 months, 2 weeks ago

Selected Answer: AC

For A - <https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/serverless-sam-cli-using-automated-tests.html#:~:text=You%20can%20use%20the%20sam%20local%20invoke%20command%20to%20manually%20test%20your%20code%20by%20running%20>

C - <https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverless-apps.html>

upvoted 1 times

 **KarBiswa** 11 months ago

Selected Answer: CD

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverless-apps.html>

upvoted 1 times

 **KarBiswa** 11 months ago

Selected Answer: AD

I will got A & D

D is nice and clear no debates

For A - <https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/serverless-sam-cli-using-automated-tests.html#:~:text=You%20can%20use%20the%20sam%20local%20invoke%20command%20to%20manually%20test%20your%20code%20by%20running%20>

upvoted 1 times

 **KarBiswa** 11 months ago

changing the option to C, D

upvoted 1 times

SerialiDr 1 year ago

Selected Answer: CD

C. Enable gradual deployments through AWS SAM templates: Gradual deployments allow you to safely deploy your application while exposing new versions to only a portion of your traffic. This approach is ideal for testing new updates in a production environment without impacting all users. AWS SAM supports the configuration of deployment preferences directly within the SAM template.

D. Set the deployment preference type to Canary10Percent30Minutes. Use hooks to test the deployment: The Canary deployment type is suitable for gradually introducing a new version of the Lambda function. In this case, "Canary10Percent30Minutes" means that 10% of the traffic will be directed to the new version for 30 minutes. If no issues are detected, the rest of the traffic is shifted to the new version. The use of hooks allows for automated tests to run against the new deployment, ensuring its stability before full traffic shift.

upvoted 2 times

c9ebec2 1 year, 1 month ago

Selected Answer: AD

A: test code during deploy <https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/serverless-sam-cli-using-automated-tests.html>

D: Canary, to deploy a small percentage

upvoted 2 times

PrakashM14 1 year, 3 months ago

Selected Answer: CD

C. Enable gradual deployments through AWS SAM templates.

D. Set the deployment preference type to Canary10Percent30Minutes. Use hooks to test the deployment.

upvoted 4 times

Digo30sp 1 year, 3 months ago

Selected Answer: CE

The correct answers are (C) and (E).

(C) is the most important step because it allows you to deploy new Lambda function updates to a small percentage of your traffic.

(E) is the second most important step because it allows you to test new Lambda function updates using hooks.

upvoted 1 times

A company is using AWS CloudFormation to deploy a two-tier application. The application will use Amazon RDS as its backend database. The company wants a solution that will randomly generate the database password during deployment. The solution also must automatically rotate the database password without requiring changes to the application.

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

- A. Use an AWS Lambda function as a CloudFormation custom resource to generate and rotate the password.
- B. Use an AWS Systems Manager Parameter Store resource with the SecureString data type to generate and rotate the password.
- C. Use a cron daemon on the application's host to generate and rotate the password.
- D. Use an AWS Secrets Manager resource to generate and rotate the password.

Correct Answer: D

Community vote distribution

D (100%)

  **dilleman** Highly Voted 9 months, 3 weeks ago

Selected Answer: D

D is correct

upvoted 6 times

  **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 2 times

  **Digo30sp** 10 months ago

Selected Answer: D

D) <https://www.examttopics.com/discussions/amazon/view/88814-exam-aws-certified-developer-associate-topic-1-question-270/>

upvoted 4 times

A developer has been asked to create an AWS Lambda function that is invoked any time updates are made to items in an Amazon DynamoDB table. The function has been created, and appropriate permissions have been added to the Lambda execution role. Amazon DynamoDB streams have been enabled for the table, but the function is still not being invoked.

Which option would enable DynamoDB table updates to invoke the Lambda function?

- A. Change the StreamViewType parameter value to NEW_AND_OLD_IMAGES for the DynamoDB table.
- B. Configure event source mapping for the Lambda function.
- C. Map an Amazon Simple Notification Service (Amazon SNS) topic to the DynamoDB streams.
- D. Increase the maximum runtime (timeout) setting of the Lambda function.

Correct Answer: B

Community vote distribution

B (100%)

🗨️ **albert_kuo** 4 months ago

Selected Answer: B

```
aws lambda create-event-source-mapping \  
--function-name MyLambdaFunction \  
--event-source-arn arn:aws:dynamodb:region:account-id:table/MyTable/stream/2023-01-01T12:00:00.000 \  
--batch-size 100 \  
--starting-position LATEST  
upvoted 1 times
```

🗨️ **65703c1** 8 months, 1 week ago

Selected Answer: B

B is the correct answer.
upvoted 2 times

🗨️ **KarBiswa** 11 months ago

Selected Answer: B

<https://docs.aws.amazon.com/lambda/latest/dg/with-ddb.html>
upvoted 2 times

🗨️ **joshnort** 12 months ago

Selected Answer: B

Configure event source mapping for the Lambda function after enabling Streams
upvoted 1 times

🗨️ **dilleman** 1 year, 3 months ago

Selected Answer: B

B is the only option that makes sense here
upvoted 3 times

🗨️ **Digo30sp** 1 year, 3 months ago

Selected Answer: B

B) <https://www.examttopics.com/discussions/amazon/view/4365-exam-aws-certified-developer-associate-topic-1-question-35/#>
upvoted 3 times

A developer needs to deploy an application running on AWS Fargate using Amazon ECS. The application has environment variables that must be passed to a container for the application to initialize.

How should the environment variables be passed to the container?

- A. Define an array that includes the environment variables under the environment parameter within the service definition.
- B. Define an array that includes the environment variables under the environment parameter within the task definition.
- C. Define an array that includes the environment variables under the entryPoint parameter within the task definition.
- D. Define an array that includes the environment variables under the entryPoint parameter within the service definition.

Correct Answer: B

Community vote distribution

B (100%)

 **Digo30sp** Highly Voted 1 year, 3 months ago

Selected Answer: B

B) <https://www.examtopycs.com/discussions/amazon/view/28795-exam-aws-certified-developer-associate-topic-1-question-108/>
upvoted 7 times

 **albert_kuo** Most Recent 4 months ago

Selected Answer: B

```
{
  "containerDefinitions": [
    {
      "name": "my-container",
      "image": "my-image",
      "environment": [
        {
          "name": "ENV_VAR_1",
          "value": "value1"
        },
        {
          "name": "ENV_VAR_2",
          "value": "value2"
        }
      ]
    }
  ]
}
```

upvoted 1 times

 **65703c1** 8 months, 1 week ago

Selected Answer: B

B is the correct answer.
upvoted 1 times

 **joshnort** 12 months ago

Selected Answer: B

<https://docs.aws.amazon.com/AmazonECS/latest/developerguide/taskdef-envfiles.html>
upvoted 3 times

 **TanTran04** 1 year, 1 month ago

Selected Answer: B

Following ChatGPT 3.5, The correct option is B
When using Amazon ECS, the task definition is where you define parameters for your containers, including environment variables. The

environment parameter within the task definition allows you to specify environment variables for your containers. This approach provides a clear separation of concerns, allowing you to define the environment variables at the task definition level, which is then used by the service when running tasks.

Option A is incorrect because the environment variables are typically defined in the task definition rather than the service definition.

Options C and D are incorrect because the `entryPoint` parameter is used for specifying the entry point (command) for the container, not for defining environment variables.

upvoted 3 times

 **dilleman** 1 year, 3 months ago

Selected Answer: B

B is correct

upvoted 4 times

A development team maintains a web application by using a single AWS RDS, template. The template defines web servers and an Amazon RDS database. The team uses the CloudFormation template to deploy the CloudFormation stack to different environments.

During a recent application deployment, a developer caused the primary development database to be dropped and recreated. The result of this incident was a loss of data. The team needs to avoid accidental database deletion in the future.

Which solutions will meet these requirements? (Choose two.)

- A. Add a CloudFormation DeletionPolicy attribute with the Retain value to the database resource.
- B. Update the CloudFormation stack policy to prevent updates to the database.
- C. Modify the database to use a Multi-AZ deployment.
- D. Create a CloudFormation stack set for the web application and database deployments.
- E. Add a CloudFormation DeletionPolicy attribute with the Retain value to the stack.

Correct Answer: AB

Community vote distribution

AB (100%)

🗳️ 👤 **65703c1** 2 months, 1 week ago

Selected Answer: AB

AB is the correct answer.

upvoted 3 times

🗳️ 👤 **joshnort** 6 months ago

Selected Answer: AB

<https://repost.aws/knowledge-center/cloudformation-accidental-updates>

upvoted 2 times

🗳️ 👤 **SerialiDr** 6 months, 1 week ago

Selected Answer: AB

A. Add a CloudFormation DeletionPolicy attribute with the Retain value to the database resource.

The DeletionPolicy attribute can be used in the CloudFormation template to protect a resource from being accidentally deleted. By setting the DeletionPolicy to Retain, the resource is retained when the stack is deleted, thus preventing accidental data loss. This should be applied directly to the database resource in the CloudFormation template.

B. Update the CloudFormation stack policy to prevent updates to the database.

A stack policy can be used to prevent changes to certain resources during stack updates. By defining a stack policy that prohibits actions on the database resource, the team can prevent accidental modifications or deletions of the database during stack updates.

upvoted 3 times

🗳️ 👤 **Gold07** 9 months, 2 weeks ago

The answer is A and D

upvoted 2 times

🗳️ 👤 **Digo30sp** 10 months ago

Selected Answer: AB

A and B) <https://www.examttopics.com/discussions/amazon/view/103521-exam-aws-certified-developer-associate-dva-c02-topic-1/#>

upvoted 4 times

A developer is storing sensitive data generated by an application in Amazon S3. The developer wants to encrypt the data at rest. A company policy requires an audit trail of when the AWS Key Management Service (AWS KMS) key was used and by whom.

Which encryption option will meet these requirements?

- A. Server-side encryption with Amazon S3 managed keys (SSE-S3)
- B. Server-side encryption with AWS KMS managed keys (SSE-KMS)
- C. Server-side encryption with customer-provided keys (SSE-C)
- D. Server-side encryption with self-managed keys

Correct Answer: B

Community vote distribution

B (100%)

🗨️ **tsangckl** 7 months, 2 weeks ago

This appear at 17 Jun exam
upvoted 1 times

🗨️ **65703c1** 8 months, 1 week ago

Selected Answer: B

B is the correct answer.
upvoted 3 times

🗨️ **SerialiDr** 1 year ago

Selected Answer: B

SSE-KMS provides an additional layer of security by requiring separate permissions for the use of an encryption key to the bucket. This option also provides an audit trail by logging the use of the key in AWS CloudTrail, which is a requirement according to the company policy. The audit trail includes information about who used the key and when it was used, which fulfills the requirement for auditing.
upvoted 2 times

🗨️ **dilleman** 1 year, 3 months ago

Selected Answer: B

B, since we need an audit trail of the AWK KMS key then this is the one to use.
upvoted 3 times

🗨️ **Digo30sp** 1 year, 3 months ago

Selected Answer: B

B) <https://www.examttopics.com/discussions/amazon/view/28801-exam-aws-certified-developer-associate-topic-1-question-217/>
upvoted 4 times

A company has an ecommerce application. To track product reviews, the company's development team uses an Amazon DynamoDB table.

Every record includes the following:

- A Review ID, a 16-digit universally unique identifier (UUID)
- A Product ID and User ID, 16-digit UUIDs that reference other tables
- A Product Rating on a scale of 1-5
- An optional comment from the user

The table partition key is the Review ID. The most performed query against the table is to find the 10 reviews with the highest rating for a given product.

Which index will provide the FASTEST response for this query?

- A. A global secondary index (GSI) with Product ID as the partition key and Product Rating as the sort key
- B. A global secondary index (GSI) with Product ID as the partition key and Review ID as the sort key
- C. A local secondary index (LSI) with Product ID as the partition key and Product Rating as the sort key
- D. A local secondary index (LSI) with Review ID as the partition key and Product ID as the sort key

Correct Answer: A

Community vote distribution

A (100%)

65703c1 2 months, 1 week ago

Selected Answer: A

A is the correct answer.
upvoted 2 times

KarBiswa 5 months ago

Selected Answer: A

Rating for a product these two are essential and already ReviewID has partitions
upvoted 2 times

SerialiDr 6 months, 1 week ago

Selected Answer: A

This GSI allows for queries that are based on Product ID, efficiently narrowing down the reviews for a specific product. The Product Rating as the sort key enables sorting the reviews by their rating, which directly supports the need to find the top 10 reviews. GSIs also support a different partition key than the base table, which is necessary in this case since the base table's partition key is Review ID.
upvoted 3 times

dilleman 9 months, 3 weeks ago

Selected Answer: A

A should be correct
upvoted 3 times

Digo30sp 10 months ago

Selected Answer: A

A) <https://www.examttopics.com/discussions/amazon/view/88995-exam-aws-certified-developer-associate-topic-1-question-362/>
upvoted 3 times

A company needs to distribute firmware updates to its customers around the world.

Which service will allow easy and secure control of the access to the downloads at the lowest cost?

- A. Use Amazon CloudFront with signed URLs for Amazon S3.
- B. Create a dedicated Amazon CloudFront Distribution for each customer.
- C. Use Amazon CloudFront with AWS Lambda@Edge.
- D. Use Amazon API Gateway and AWS Lambda to control access to an S3 bucket.

Correct Answer: A

Community vote distribution

A (100%)

🗨️ **TanTran04** Highly Voted 👍 7 months, 3 weeks ago

Selected Answer: A

A. Use Amazon CloudFront with signed URLs for Amazon S3.

Using Amazon CloudFront with signed URLs is a secure and cost-effective way to control access to downloads. With signed URLs, you can generate URLs with limited time validity, ensuring that only users with the correct URL and during the specified time window can access the firmware updates. This provides both security and control over access.

Option B (Create a dedicated Amazon CloudFront Distribution for each customer) may result in higher costs and increased complexity. Option C (Use Amazon CloudFront with AWS Lambda@Edge) is more focused on customization and additional processing at the edge locations, which may not be necessary for simple access control. Option D (Use Amazon API Gateway and AWS Lambda) is more suited for managing APIs and might be an overkill for a straightforward firmware update distribution scenario.

upvoted 5 times

🗨️ **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 2 times

🗨️ **SerialIDr** 6 months, 1 week ago

Selected Answer: A

Amazon CloudFront is a content delivery network (CDN) service that can efficiently distribute files globally. Using signed URLs provides a secure method to control access to the firmware updates. Only users with valid signed URLs can download the files, ensuring controlled access. This approach is cost-effective as it leverages CloudFront's caching capabilities to reduce load on the origin (Amazon S3) and reduces data transfer costs.

upvoted 3 times

🗨️ **walala97** 6 months, 4 weeks ago

Selected Answer: A

option B, for each customer, will bring high costs

upvoted 2 times

🗨️ **dilleman** 9 months, 3 weeks ago

Selected Answer: A

A is correct

upvoted 3 times

🗨️ **Digo30sp** 10 months ago

Selected Answer: A

A) <https://www.examttopics.com/discussions/amazon/view/8792-exam-aws-certified-developer-associate-topic-1-question-179/#>

upvoted 4 times

A developer is testing an application that invokes an AWS Lambda function asynchronously. During the testing phase, the Lambda function fails to process after two retries.

How can the developer troubleshoot the failure?

- A. Configure AWS CloudTrail logging to investigate the invocation failures.
- B. Configure Dead Letter Queues by sending events to Amazon SQS for investigation.
- C. Configure Amazon Simple Workflow Service to process any direct unprocessed events.
- D. Configure AWS Config to process any direct unprocessed events.

Correct Answer: B

Community vote distribution

B (100%)

🗨️ **TanTran04** Highly Voted 7 months, 3 weeks ago

Selected Answer: B

In AWS Lambda, you can use Dead Letter Queues (DLQ) to capture and retain events that couldn't be processed successfully after a specified number of retries. By configuring a DLQ, the failed events are sent to an Amazon SQS queue, allowing you to investigate and analyze the reasons for the failures.

upvoted 6 times

🗨️ **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

🗨️ **dilleman** 9 months, 3 weeks ago

Selected Answer: B

Dead Letter Queues (DLQ) can be configured for Lambda functions to capture failed asynchronous invocations. Events that cannot be processed will be sent to an SQS queue (or an SNS topic) you specify, allowing for further investigation and reprocessing.

upvoted 4 times

🗨️ **Digo30sp** 10 months ago

Selected Answer: B

B) <https://www.examttopics.com/discussions/amazon/view/28638-exam-aws-certified-developer-associate-topic-1-question-317/#>

upvoted 2 times

A company is migrating its PostgreSQL database into the AWS Cloud. The company wants to use a database that will secure and regularly rotate database credentials. The company wants a solution that does not require additional programming overhead.

Which solution will meet these requirements?

- A. Use Amazon Aurora PostgreSQL for the database. Store the database credentials in AWS Systems Manager Parameter Store. Turn on rotation.
- B. Use Amazon Aurora PostgreSQL for the database. Store the database credentials in AWS Secrets Manager. Turn on rotation.
- C. Use Amazon DynamoDB for the database. Store the database credentials in AWS Systems Manager Parameter Store. Turn on rotation.
- D. Use Amazon DynamoDB for the database. Store the database credentials in AWS Secrets Manager. Turn on rotation.

Correct Answer: B

Community vote distribution

B (100%)

 **Digo30sp** Highly Voted 10 months ago

Selected Answer: B

B) The correct answer is (B).

Solution (B) is the best option because it meets all the requirements:

Using a database that secures and regularly changes database credentials: Amazon Aurora PostgreSQL offers built-in credential rotation, which allows you to change database credentials at regular intervals.

Solution that requires no additional programming overhead: Amazon Aurora PostgreSQL credential rotation is fully automated, so it requires no additional programming overhead.

upvoted 8 times

 **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

 **TanTran04** 7 months, 3 weeks ago

Selected Answer: B

AWS Secrets Manager is a service designed to rotate, manage, and retrieve database credentials, API keys, and other secrets. In this scenario, storing the database credentials in AWS Secrets Manager and enabling rotation ensures that your database credentials are regularly rotated without requiring additional programming overhead.

Option B is the most suitable because it specifically addresses the need for securing and regularly rotating database credentials in the AWS Cloud. Amazon Aurora PostgreSQL is a fully managed relational database service, and AWS Secrets Manager seamlessly integrates with it for secure credential management.

upvoted 1 times

A developer is creating a mobile application that will not require users to log in.

What is the MOST efficient method to grant users access to AWS resources?

- A. Use an identity provider to securely authenticate with the application.
- B. Create an AWS Lambda function to create an IAM user when a user accesses the application.
- C. Create credentials using AWS KMS and apply these credentials to users when using the application.
- D. Use Amazon Cognito to associate unauthenticated users with an IAM role that has limited access to resources.

Correct Answer: D

Community vote distribution

D (100%)

🗨️ 👤 **Digo30sp** Highly Voted 👍 10 months ago

Selected Answer: D

D) <https://www.examtopycs.com/discussions/amazon/view/4245-exam-aws-certified-developer-associate-topic-1-question-79/>
upvoted 8 times

🗨️ 👤 **65703c1** Most Recent 🕒 2 months, 1 week ago

Selected Answer: D

D is the correct answer.
upvoted 1 times

🗨️ 👤 **SerialiDr** 6 months, 1 week ago

Selected Answer: D

Amazon Cognito allows you to create unique identities for users of your application and assign permissions to these identities using IAM roles. By using Cognito's unauthenticated identities (also known as guest users), you can grant limited AWS resource access to users without requiring them to log in. This approach is secure, scalable, and does not require managing user credentials.
upvoted 3 times

🗨️ 👤 **TanTran04** 7 months, 3 weeks ago

Selected Answer: D

Amazon Cognito is designed to handle user identity and access management for mobile and web applications
upvoted 1 times

A company has developed a new serverless application using AWS Lambda functions that will be deployed using the AWS Serverless Application Model (AWS SAM) CLI.

Which step should the developer complete prior to deploying the application?

- A. Compress the application to a .zip file and upload it into AWS Lambda.
- B. Test the new AWS Lambda function by first tracing it in AWS X-Ray.
- C. Bundle the serverless application using a SAM package.
- D. Create the application environment using the `eb create my-env` command.

Correct Answer: C

Community vote distribution

C (100%)

🗨️ **albert_kuo** 4 months ago

Selected Answer: C

`sam package --template-file template.yaml --s3-bucket my-deployment-bucket --output-template-file packaged.yaml`
upvoted 1 times

🗨️ **65703c1** 8 months, 1 week ago

Selected Answer: C

C is the correct answer.
upvoted 2 times

🗨️ **SerialiDr** 1 year ago

Selected Answer: C

Using `sam package` (or `aws cloudformation package`) command, the developer can package the Lambda functions, along with any dependencies and resources defined in the SAM template, into a deployment package. This command uploads local artifacts (like Lambda function code and Swagger files) to an S3 bucket and produces a modified SAM template file, formatted for deployment.
upvoted 4 times

🗨️ **TanTran04** 1 year, 1 month ago

Selected Answer: C

C. Bundle the serverless application using a SAM package.

Before deploying a serverless application using the AWS Serverless Application Model (AWS SAM) CLI, the developer should bundle the application using the `sam package` command. This command packages and uploads the local artifacts of your serverless application to Amazon S3, and it produces a packaged AWS SAM template file that you can deploy with the `sam deploy` command.

upvoted 1 times

🗨️ **dilleman** 1 year, 3 months ago

Selected Answer: C

C is correct
upvoted 2 times

🗨️ **Digo30sp** 1 year, 3 months ago

Selected Answer: C

C) <https://www.examtopycs.com/discussions/amazon/view/28650-exam-aws-certified-developer-associate-topic-1-question-312/>
upvoted 4 times

A company wants to automate part of its deployment process. A developer needs to automate the process of checking for and deleting unused resources that supported previously deployed stacks but that are no longer used.

The company has a central application that uses the AWS Cloud Development Kit (AWS CDK) to manage all deployment stacks. The stacks are spread out across multiple accounts. The developer's solution must integrate as seamlessly as possible within the current deployment process.

Which solution will meet these requirements with the LEAST amount of configuration?

- A. In the central AWS CDK application, write a handler function in the code that uses AWS SDK calls to check for and delete unused resources. Create an AWS CloudFormation template from a JSON file. Use the template to attach the function code to an AWS Lambda function and to invoke the Lambda function when the deployment stack runs.
- B. In the central AWS CDK application, write a handler function in the code that uses AWS SDK calls to check for and delete unused resources. Create an AWS CDK custom resource. Use the custom resource to attach the function code to an AWS Lambda function and to invoke the Lambda function when the deployment stack runs.
- C. In the central AWS CDK, write a handler function in the code that uses AWS SDK calls to check for and delete unused resources. Create an API in AWS Amplify. Use the API to attach the function code to an AWS Lambda function and to invoke the Lambda function when the deployment stack runs.
- D. In the AWS Lambda console, write a handler function in the code that uses AWS SDK calls to check for and delete unused resources. Create an AWS CDK custom resource. Use the custom resource to import the Lambda function into the stack and to invoke the Lambda function when the deployment stack runs.

Correct Answer: B

Community vote distribution

B (100%)

🗨️ **tsangckl** 7 months, 2 weeks ago

This appear at 17 Jun exam
upvoted 1 times

🗨️ **65703c1** 8 months, 1 week ago

Selected Answer: B

B is the correct answer.
upvoted 2 times

🗨️ **SerialIDr** 1 year ago

Selected Answer: B

This approach integrates directly into the existing AWS CDK deployment process. By writing a handler function within the CDK application, the developer can leverage the AWS SDK to programmatically identify and delete unused resources. The AWS CDK custom resource can then be used to invoke this function as part of the deployment process. This solution is efficient as it keeps everything within the CDK ecosystem and minimizes additional external configurations.

upvoted 3 times

🗨️ **dilleman** 1 year, 3 months ago

Selected Answer: B

B. In the central AWS CDK application, write a handler function in the code that uses AWS SDK calls to check for and delete unused resources. Create an AWS CDK custom resource. Use the custom resource to attach the function code to an AWS Lambda function and to invoke the Lambda function when the deployment stack runs.

upvoted 1 times

🗨️ **Digo30sp** 1 year, 3 months ago

Selected Answer: B

The correct answer is (B).

Solution (B) is the best option because:

Requires the LEAST amount of configuration: Solution (B) uses an AWS CDK custom resource, which is a type of resource that can be defined in AWS CDK code. Custom resources are a convenient way to add custom functionality to your AWS CloudFormation stacks.

Integrates seamlessly into the current deployment process: Solution (B) uses the AWS CDK custom resource to attach function code to an AWS Lambda function and to invoke the Lambda function when the deployment stack runs. This means that the solution does not require any changes to the existing AWS CDK code.

upvoted 4 times

A company built a new application in the AWS Cloud. The company automated the bootstrapping of new resources with an Auto Scaling group by using AWS CloudFormation templates. The bootstrap scripts contain sensitive data.

The company needs a solution that is integrated with CloudFormation to manage the sensitive data in the bootstrap scripts.

Which solution will meet these requirements in the MOST secure way?

- A. Put the sensitive data into a CloudFormation parameter. Encrypt the CloudFormation templates by using an AWS Key Management Service (AWS KMS) key.
- B. Put the sensitive data into an Amazon S3 bucket. Update the CloudFormation templates to download the object from Amazon S3 during bootstrap.
- C. Put the sensitive data into AWS Systems Manager Parameter Store as a secure string parameter. Update the CloudFormation templates to use dynamic references to specify template values.
- D. Put the sensitive data into Amazon Elastic File System (Amazon EFS). Enforce EFS encryption after file system creation. Update the CloudFormation templates to retrieve data from Amazon EFS.

Correct Answer: C

Community vote distribution

C (94%) 6%

🗨️ **Digo30sp** Highly Voted 10 months ago

Selected Answer: C

The correct answer is (C).

Solution (C) is the best option because:

It's the most secure solution: Sensitive data is stored in AWS Systems Manager Parameter Store, which is a secret management service managed by AWS. Secure string parameters in AWS Systems Manager Parameter Store are encrypted with an AWS KMS key.

It's integrated with CloudFormation: Secure string parameters can be referenced in CloudFormation templates using dynamic references. This means that sensitive data does not need to be stored in CloudFormation code.

upvoted 7 times

🗨️ **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗨️ **KarBiswa** 5 months ago

Selected Answer: C

<https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-parameter-store.html>

upvoted 1 times

🗨️ **TanTran04** 7 months, 3 weeks ago

Selected Answer: C

AWS Systems Manager Parameter Store is a secure and scalable solution for storing configuration data, including sensitive information. In this case, using a secure string parameter allows you to store the sensitive data in Parameter Store in an encrypted form.

Option C is the most secure because it leverages AWS Systems Manager Parameter Store's capabilities for securely storing sensitive data, and dynamic references allow you to directly reference the parameter values in CloudFormation templates. This approach avoids exposing sensitive data in the templates themselves and provides a central and secure storage solution for sensitive configuration information.

upvoted 2 times

🗨️ **kashtelyan** 9 months, 2 weeks ago

Selected Answer: A



A option leverages CloudFormation parameters, which can securely store sensitive data. By using an AWS KMS key to encrypt the CloudFormation templates, you ensure that the sensitive data is protected. It follows the principle of least privilege and provides secure access to sensitive information directly within CloudFormation.

Option B is less secure because it involves storing sensitive data in an S3 bucket, which could be compromised.

Option C suggests using AWS Systems Manager Parameter Store, which is secure, but using CloudFormation parameters and KMS keys provides an integrated solution directly within CloudFormation.

Option D involves Amazon EFS, which is typically used for file storage and is not designed for securely storing sensitive data directly within CloudFormation.

upvoted 1 times

  **dilleman** 9 months, 3 weeks ago

Selected Answer: C

C is the correct choice. Parameter Store's secure string parameter encrypts the data using AWS KMS

upvoted 4 times

A company needs to set up secure database credentials for all its AWS Cloud resources. The company's resources include Amazon RDS DB instances, Amazon DocumentDB clusters, and Amazon Aurora DB instances. The company's security policy mandates that database credentials be encrypted at rest and rotated at a regular interval.

Which solution will meet these requirements MOST securely?

- A. Set up IAM database authentication for token-based access. Generate user tokens to provide centralized access to RDS DB instances, Amazon DocumentDB clusters, and Aurora DB instances.
- B. Create parameters for the database credentials in AWS Systems Manager Parameter Store. Set the Type parameter to SecureString. Set up automatic rotation on the parameters.
- C. Store the database access credentials as an encrypted Amazon S3 object in an S3 bucket. Block all public access on the S3 bucket. Use S3 server-side encryption to set up automatic rotation on the encryption key.
- D. Create an AWS Lambda function by using the SecretsManagerRotationTemplate template in the AWS Secrets Manager console. Create secrets for the database credentials in Secrets Manager. Set up secrets rotation on a schedule.

Correct Answer: D

Community vote distribution

D (100%)

wh1t4k3r 5 months ago

I love how they added the lambda b**is**t just to confuse who's doing the test.
upvoted 1 times

65703c1 8 months, 1 week ago

Selected Answer: D

D is the correct answer.
upvoted 2 times

TanTran04 1 year, 1 month ago

Selected Answer: D

Rotate auto SecretsManager
upvoted 3 times

nickolaj 1 year, 3 months ago

<https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-manager/>
upvoted 2 times

dilleman 1 year, 3 months ago

Selected Answer: D

the best and most secure option is:
D. Create an AWS Lambda function by using the SecretsManagerRotationTemplate template in the AWS Secrets Manager console.
upvoted 4 times

Digo30sp 1 year, 3 months ago

Selected Answer: D

The correct answer is (D).

Solution (D) is the best option because:

It's the most secure solution: AWS Secrets Manager is an AWS-managed secrets management service that provides encryption at rest and automatic secret rotation.

Meets the company's security requirements: The solution meets the company's security requirements because:

Database credentials are encrypted at rest using AWS Key Management Service (AWS KMS).

Database credentials are automatically rotated at regular intervals.

upvoted 4 times

  **fordiscussiontwo** 1 year, 3 months ago

DDDDDD

upvoted 3 times

A developer has created an AWS Lambda function that makes queries to an Amazon Aurora MySQL DB instance. When the developer performs a test, the DB instance shows an error for too many connections.

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

- A. Create a read replica for the DB instance. Query the replica DB instance instead of the primary DB instance.
- B. Migrate the data to an Amazon DynamoDB database.
- C. Configure the Amazon Aurora MySQL DB instance for Multi-AZ deployment.
- D. Create a proxy in Amazon RDS Proxy. Query the proxy instead of the DB instance.

Correct Answer: D

Community vote distribution

D (100%)

🗉 **jasonkym** 2 months, 1 week ago

Please educated me why is it not A?

upvoted 1 times

🗉 **65703c1** 8 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 2 times

🗉 **Dzok5050** 8 months, 2 weeks ago

Selected Answer: D

Too many Connections = Proxy

upvoted 2 times

🗉 **SathyaJS** 10 months ago

Selected Answer : D

upvoted 1 times

🗉 **TanTran04** 1 year, 1 month ago

too many connections => proxy

upvoted 1 times

🗉 **dilleman** 1 year, 3 months ago

Selected Answer: D

D.

RDS Proxy sits between the application and the database to manage and pool connections, reducing the chance of exhausting database connections when many Lambda functions try to connect simultaneously.

upvoted 3 times

🗉 **Digo30sp** 1 year, 3 months ago

Selected Answer: D

D) <https://www.examtopycs.com/discussions/amazon/view/88969-exam-aws-certified-developer-associate-topic-1-question-358/>

upvoted 1 times

🗉 **fordiscussiontwo** 1 year, 3 months ago

DDDDDDDDDD

upvoted 3 times

A developer is creating a new REST API by using Amazon API Gateway and AWS Lambda. The development team tests the API and validates responses for the known use cases before deploying the API to the production environment.

The developer wants to make the REST API available for testing by using API Gateway locally.

Which AWS Serverless Application Model Command Line Interface (AWS SAM CLI) subcommand will meet these requirements?

- A. Sam local invoke
- B. Sam local generate-event
- C. Sam local start-lambda
- D. Sam local start-api

Correct Answer: D

Community vote distribution

D (100%)

  **Digo30sp** Highly Voted 10 months ago

Selected Answer: D

The correct answer is (D).

The AWS SAM CLI sam local start-api subcommand is used to start a local API Gateway instance. This allows you to test your REST API locally before deploying it to the production environment.

The other subcommands will not meet the developer's requirements:

Local invocation of Sam is used to invoke a Lambda function locally.

Sam's local event generation is used to generate a local event file to be used to invoke a Lambda function locally.

Sam local start-lambda is used to start a local instance of a Lambda function.


upvoted 9 times

  **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: D

D is the correct answer.



upvoted 1 times

  **KarBiswa** 5 months ago

Selected Answer: D

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/sam-cli-command-reference-sam-local-start-api.html>

upvoted 1 times

  **dilleman** 9 months, 3 weeks ago

Selected Answer: D

D is correct

upvoted 3 times

  **fordiscussiontwo** 10 months ago

DDDDDDDDDD

upvoted 4 times


A company has a serverless application on AWS that uses a fleet of AWS Lambda functions that have aliases. The company regularly publishes new Lambda function by using an in-house deployment solution. The company wants to improve the release process and to use traffic shifting. A newly published function version should initially make available only to a fixed percentage of production users.

Which solution will meet these requirements?

- A. Configure routing on the alias of the new function by using a weighted alias.
- B. Configure a canary deployment type for Lambda.
- C. Configure routing on the new versions by using environment variables.
- D. Configure a linear deployment type for Lambda.

Correct Answer: A

Community vote distribution



Digo30sp Highly Voted 1 year, 3 months ago

Selected Answer: A

The correct answer is (A).

Weighted aliases allow you to route traffic to different versions of a function based on weights that you assign. This allows you to implement a canary deployment, where you initially route a small percentage of your traffic to the new version of the function, and then gradually increase the percentage as you gain confidence in the new version.

upvoted 8 times

rimaSamir 11 months, 2 weeks ago

If we need Canary deployment, then why not B ?

How you will use A in automated deployment?

upvoted 3 times

albert_kuo 4 months ago

While canary deployments are a valid strategy for gradual rollouts, Lambda doesn't have a built-in "canary deployment type." This option is misleading and not applicable in the context of Lambda functions.

upvoted 1 times

albert_kuo 6 months, 1 week ago

the key word is "fixed percentage"

upvoted 2 times

65703c1 Most Recent 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

DeaconStJohn 10 months, 2 weeks ago

Selected Answer: B

I am struggling to see how the correct answer isn't canary. please feel free to enlighten me as I am at a loss how this question description is anything but canary

upvoted 1 times

DeaconStJohn 10 months, 2 weeks ago

is this question a case of what naming convention is used within Lambda service. i.e. Canary deployments via Weighted Aliases.

upvoted 1 times

DeaconStJohn 10 months, 2 weeks ago

Still very on the fence with this one.

My key take aways are that the question says an in house deployment solution and not "codedeploy."

By using weighted aliases we are in fact performing a canary deployments.

Bit of a question.

upvoted 1 times

  **ahadh7621** 6 months, 2 weeks ago

My thought process is this. They can't use Canary Deployment because that is specifically for AWS CodeDeploy:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-configurations.html>

They are using an in-house deployment method, so AWS canary deployments aren't applicable. They can, however, use routing configuration on an alias to send a portion of traffic to a second function version. For example, you can reduce the risk of deploying a new version by configuring the alias to send most of the traffic to the existing version, and only a small percentage of traffic to the new version.

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html#configuring-alias-routing>

upvoted 2 times



  **KarBiswa** 11 months ago

Selected Answer: A

[https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html#configuring-alias-](https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html#configuring-alias-routing)

[routing:~:text=function%20version.%20For%20example%2C%20you%20can%20reduce%20the%20risk%20of%20deploying%20a%20new%20version%20by%20](https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html#configuring-alias-routing)

upvoted 2 times

  **NijeshT** 1 year, 2 months ago

Answer is A.

weighted aliases offer fixed, predefined percentages

upvoted 4 times

  **fordiscussiontwo** 1 year, 3 months ago

AAAAAAAAAAAA

upvoted 4 times

A company has an application that stores data in Amazon RDS instances. The application periodically experiences surges of high traffic that cause performance problems. During periods of peak traffic, a developer notices a reduction in query speed in all database queries.

The team's technical lead determines that a multi-threaded and scalable caching solution should be used to offload the heavy read traffic. The solution needs to improve performance.

Which solution will meet these requirements with the LEAST complexity?

- A. Use Amazon ElastiCache for Memcached to offload read requests from the main database.
- B. Replicate the data to Amazon DynamoDB set up a DynamoDB Accelerator (DAX) cluster.
- C. Configure the Amazon RDS instances to use Multi-AZ deployment with one standby instance. Offload read requests from the main database to the standby instance.
- D. Use Amazon ElastiCache for Redis to offload read requests from the main database.

Correct Answer: A

Community vote distribution

A (100%)

 **kashtelyan** Highly Voted 1 year, 3 months ago

Selected Answer: A

When deciding between Memcached and Redis, here are a few questions to consider:

Is object caching your primary goal, for example to offload your database? If so, use Memcached.

<https://docs.aws.amazon.com/whitepapers/latest/scale-performance-elasticache/memcached-vs.-redis.html>


upvoted 8 times

 **albert_kuo** Most Recent 4 months ago

Selected Answer: A

option A (using ElastiCache for Memcached) provides the best balance of meeting the requirements (multi-threaded, scalable caching to improve performance) while maintaining the least complexity.

upvoted 1 times

 **65703c1** 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

 **tsdsmth** 12 months ago

Selected Answer: A

A. If you're looking for a multi-threaded solution, then ElastiCache for Memcached (not Redis) is the solution.

upvoted 1 times

 **SerialiDr** 1 year ago

Selected Answer: A

A. Use Amazon ElastiCache for Memcached to offload read requests from the main database.

ElastiCache for Memcached is a good fit for this scenario. It's a high-performance, distributed, in-memory caching system that can easily scale to manage surges in read traffic. It's simple to set up and integrate with an existing RDS instance.

D. Use Amazon ElastiCache for Redis to offload read requests from the main database.

ElastiCache for Redis also offers high performance and is capable of handling surges in read traffic. Redis provides more advanced data structures and features compared to Memcached, like persistence, built-in replication, and support for complex data types. However, it might be more complex to set up and manage than Memcached, depending on the use case.

upvoted 3 times

🗨️ 👤 **Digo30sp** 1 year, 3 months ago

Selected Answer: A

The correct answer is (A).

Amazon ElastiCache for Memcached is a scalable, multithreaded caching solution that can be used to offload heavy read traffic from Amazon RDS instances. ElastiCache for Memcached is easy to configure and manage, making it a low-effort solution to meet technical lead requirements.
upvoted 3 times

🗨️ 👤 **fordiscussiontwo** 1 year, 3 months ago

AAAAAAAAA

upvoted 3 times

A developer must provide an API key to an AWS Lambda function to authenticate with a third-party system. The Lambda function will run on a schedule. The developer needs to ensure that the API key remains encrypted at rest.

Which solution will meet these requirements?

- A. Store the API key as a Lambda environment variable by using an AWS Key Management Service (AWS KMS) customer managed key.
- B. Configure the application to prompt the user to provide the password to the Lambda function on the first run.
- C. Store the API key as a value in the application code.
- D. Use Lambda@Edge and only communicate over the HTTPS protocol.

Correct Answer: A

Community vote distribution

A (100%)

 **Digo30sp** Highly Voted 9 months, 4 weeks ago

Selected Answer: A

The correct answer is (A).

Storing the API key as a Lambda environment variable using an AWS Key Management Service (AWS KMS) customer-managed key is the most secure solution. AWS KMS is a managed encryption service that provides customer-managed keys. Customer-managed keys are encrypted with an AWS KMS master key, which is stored in an AWS KMS vault.


upvoted 8 times

 **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

 **SerialiDr** 6 months, 1 week ago

Selected Answer: A

Lambda environment variables can be encrypted using a customer managed key in AWS KMS. This approach ensures that the API key is encrypted at rest and seamlessly integrated into the Lambda function. When the function is executed, it can access the decrypted value of the API key for authenticating with the third-party system.

upvoted 3 times

 **fordiscussiontwo** 10 months ago

AAAAAAAAAA

upvoted 2 times


An IT department uses Amazon S3 to store sensitive images. After more than 1 year, the company moves the images into archival storage. The company rarely accesses the images, but the company wants a storage solution that maximizes resiliency. The IT department needs access to the images that have been moved to archival storage within 24 hours.

Which solution will meet these requirements MOST cost-effectively?

- A. Use S3 Standard-Infrequent Access (S3 Standard-IA) to store the images. Use S3 Glacier Deep Archive with standard retrieval to store and retrieve archived images.
- B. Use S3 Standard-Infrequent Access (S3 Standard-IA) to store the images. Use S3 Glacier Deep Archive with bulk retrieval to store and retrieve archived images.
- C. Use S3 Intelligent-Tiering to store the images. Use S3 Glacier Deep Archive with standard retrieval to store and retrieve archived images.
- D. Use S3 One Zone-Infrequent Access (S3 One Zone-IA) to store the images. Use S3 Glacier Deep Archive with bulk retrieval to store and retrieve archived images.

Correct Answer: A

Community vote distribution



Learning4life Highly Voted 1 year, 3 months ago

A is correct. The requirement of maximizing resiliency rules out One Zone. Standard recover is within 12 hours, which fits the requirement of within 24 hours. <https://docs.aws.amazon.com/AmazonS3/latest/userguide/restoring-objects-retrieval-options.html>
upvoted 8 times

f271c23 Most Recent 1 month, 3 weeks ago

Selected Answer: A

I had incorrectly selected B but the right option is A. the main reason is the retrieval time in the standard way is within 12 hours which meets the requirement
upvoted 1 times

MasoudK 3 months, 3 weeks ago

Option A is not most Cost effective the standard approach maximize resiliency but is more expensive than Option B. Option B handles both goals.
upvoted 1 times

albert_kuo 4 months ago

Selected Answer: B

The average retrieval time for bulk retrievals typically ranges from 5 to 12 hours, but it can take up to 48 hours. This method is suitable for cold storage data, making it a highly cost-effective option when retrieval frequency is low, and there are no stringent requirements on retrieval time.
upvoted 2 times

lambdaFun 3 months, 3 weeks ago

bulk retrievals take 48 hours

standard retrievals take 24hours

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/restoring-objects-retrieval-options.html>

Refer the table 2 row

upvoted 1 times

Saurabh04 5 months, 2 weeks ago

Selected Answer: B

MOST cost-effective solution is:

B. Use S3 Standard-Infrequent Access (S3 Standard-IA) to store the images. Use S3 Glacier Deep Archive with bulk retrieval to store and retrieve archived images.

Here's why:



S3 Standard-Infrequent Access (S3 Standard-IA):

Provides a balance between cost and retrieval speed.

Suitable for long-lived, less frequently accessed data.

Accessible within hours.



upvoted 2 times

  **Anandesh** 6 months, 2 weeks ago

Selected Answer: A

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/glacier-storage-classes.html#archival-storage>

upvoted 1 times

  **65703c1** 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

  **KarBiswa** 11 months ago

Selected Answer: C

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/restoring-objects-retrieval-options.html#:~:text=Deep%20Archive%20or-,S3%20Intelligent%2DTiering%20Deep%20Archive%20Access,-Not%20available>

upvoted 1 times

  **DeaconStJohn** 10 months, 1 week ago

Initially I thought C also. However, lifecycle policies seem to be better for this use case. S3-IT will start at standard pricing, after 30 days > IA, after 90 days > archive instant retrieval. None of these are the most cost effective. S3-IT works well for use cases where there is no defined policy in place, i.e. after 1 year move to archive. reqs state archive after 365 days. s3-IT will action this after 90 days depending on access patterns.

upvoted 1 times

  **SerialiDr** 1 year ago

Selected Answer: A

S3 Standard-IA is designed for data that is accessed less frequently but requires rapid access when needed. It offers a lower storage cost while still providing high durability, availability, and performance.

S3 Glacier Deep Archive is the most cost-effective option for archival storage in AWS and is designed for data that is accessed very rarely. The standard retrieval option in Glacier Deep Archive typically returns data within 12 hours, meeting the requirement of access within 24 hours.

upvoted 3 times

  **_YaWeb** 1 year ago

ChatGPT goes with B

upvoted 2 times

  **DeaconStJohn** 10 months, 1 week ago

As a society we need to learn to challenge AI models.

https://aws.amazon.com/s3/faqs/#Amazon_S3_Glacier_Deep_Archive

When restoring an archived object, you can specify one of the following options in the Tier element of the request body: Standard is the default tier and lets you access any of your archived objects within 12 hours, with retrievals typically starting within 9 hours when initiated using S3 Batch Operations. Bulk lets you retrieve large amounts of data, even petabytes of data, inexpensively and typically completes within 48 hours.

upvoted 2 times

  **dostonbekabdullaev** 1 year ago

Selected Answer: A

AAAAAAAAA

upvoted 1 times

  **Certified101** 1 year, 1 month ago

Selected Answer: A

A is correct -Bulk retrieval is 48hours

upvoted 1 times

  **TanTran04** 1 year, 1 month ago

Selected Answer: B

With Option A: Standard retrieval would provide faster access to the archived images (typically within 3-5 hours), it is more expensive than Bulk retrieval. Since the company has indicated they can wait up to 24 hours for access, the slower but cheaper

=> Option B is the best choice.

upvoted 1 times

🗨️ **Hanny** 1 year, 1 month ago

Selected Answer: C

C. Use S3 Intelligent-Tiering to store the images. Use S3 Glacier Deep Archive with standard retrieval to store and retrieve archived images.

upvoted 2 times

🗨️ **tqiu654** 1 year, 1 month ago

Selected Answer: B

GPT: B. Use S3 Standard-Infrequent Access (S3 Standard-IA) to store the images. Use S3 Glacier Deep Archive and select Batch Retrieval to store and retrieve archived images.

upvoted 2 times

🗨️ **hcsaba1982** 1 year, 3 months ago

Selected Answer: C

A : Glacier Deep Archive is cheaper than Standard-IA.

C : Standard archival is 12h.

B : bulk retrieval is 48h

D : S3 One Zone-IA - cross-out due to "maximizes resiliency"

upvoted 4 times

🗨️ **ut18** 1 year, 3 months ago

Check the requirement :

The IT department needs access to the images that have been moved to archival storage within 24 hours.

upvoted 1 times

🗨️ **Cerakoted** 1 year, 3 months ago

Selected Answer: A

It is A

upvoted 2 times

A developer is building a serverless application by using the AWS Serverless Application Model (AWS SAM). The developer is currently testing the application in a development environment. When the application is nearly finished, the developer will need to set up additional testing and staging environments for a quality assurance team.


The developer wants to use a feature of the AWS SAM to set up deployments to multiple environments.

Which solution will meet these requirements with the LEAST development effort?

- A. Add a configuration file in TOML format to group configuration entries to every environment. Add a table for each testing and staging environment. Deploy updates to the environments by using the `sam deploy` command and the `--config-env` flag that corresponds to each environment.
- B. Create additional AWS SAM templates for each testing and staging environment. Write a custom shell script that uses the `sam deploy` command and the `--template-file` flag to deploy updates to the environments.
- C. Create one AWS SAM configuration file that has default parameters. Perform updates to the testing and staging environments by using the `--parameter-overrides` flag in the AWS SAM CLI and the parameters that the updates will override.
- D. Use the existing AWS SAM template. Add additional parameters to configure specific attributes for the serverless function and database table resources that are in each environment. Deploy updates to the testing and staging environments by using the `sam deploy` command.

Correct Answer: A

Community vote distribution



SerialIDr Highly Voted 1 year ago

Selected Answer: A

To set up deployments to multiple environments with the least development effort in a serverless application using the AWS Serverless Application Model (AWS SAM), the developer can utilize a configuration file in TOML format with grouped configuration entries for each environment. This approach allows for easy management of different environment configurations and streamlines the deployment process. The specific steps would include:

Creating a configuration file in TOML format: This file will include a table for each testing and staging environment, where each table contains the specific configuration for that environment.

Using the `sam deploy` command with the `--config-env` flag: This flag allows specifying which environment configuration to use for the deployment, corresponding to the tables defined in the configuration file.

This solution aligns with Option A:

upvoted 9 times

Jing2023 Highly Voted 1 year, 3 months ago

Selected Answer: A

A should be correct

reference this stackoverflow post <https://stackoverflow.com/questions/68826108/how-to-deploy-to-different-environments-with-aws-sam>

upvoted 8 times

albert_kuo Most Recent 4 months ago

Selected Answer: A

`sam deploy --config-env default`

`sam deploy --config-env testing`

`sam deploy --config-env production`

upvoted 1 times

KennethNg923 5 months, 1 week ago

Selected Answer: A

AWS SAM supports configuration files in TOML format, which allows you to define multiple environments in a single file.

upvoted 1 times

Anandesh 6 months, 3 weeks ago

Selected Answer: C

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/serverless-sam-cli-config.html#serverless-sam-cli-config-default>

upvoted 1 times

🗨️ 👤 **65703c1** 8 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

🗨️ 👤 **41eb566** 10 months, 2 weeks ago

Selected Answer: C

C. Create one AWS SAM configuration file that has default parameters. Perform updates to the testing and staging environments by using the --parameter-overrides flag in the AWS SAM CLI and the parameters that the updates will override.

upvoted 2 times

🗨️ 👤 **KarBiswa** 11 months ago

Selected Answer: A

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/serverless-sam-cli-config.html>

upvoted 2 times

🗨️ 👤 **Certified101** 1 year, 1 month ago

Selected Answer: C

C with least development overhead

upvoted 1 times

🗨️ 👤 **TanTran04** 1 year, 1 month ago

Selected Answer: C

With at LEAST development effort, Option C is better than A

While this approach may work, it introduces additional complexity with the need for a separate configuration file, and it may not be as straightforward as using parameter overrides, as suggested in option C. The use of TOML format might be more suited for certain scenarios, but in the context of AWS SAM, which commonly relies on YAML or JSON configurations, it might be an extra layer of complexity that isn't necessary.

Option C, on the other hand, recommends using a single AWS SAM configuration file with default parameters and updating testing and staging environments using the --parameter-overrides flag. This approach is more aligned with typical AWS SAM practices and is simpler and more straightforward than managing multiple configuration files.

upvoted 3 times

🗨️ 👤 **TanTran04** 1 year, 1 month ago

With at LEAST development effort, Option C is better than A

While this approach may work, it introduces additional complexity with the need for a separate configuration file, and it may not be as straightforward as using parameter overrides, as suggested in option C. The use of TOML format might be more suited for certain scenarios, but in the context of AWS SAM, which commonly relies on YAML or JSON configurations, it might be an extra layer of complexity that isn't necessary.

Option C, on the other hand, recommends using a single AWS SAM configuration file with default parameters and updating testing and staging environments using the --parameter-overrides flag. This approach is more aligned with typical AWS SAM practices and is simpler and more straightforward than managing multiple configuration files.

upvoted 1 times

🗨️ 👤 **Hanny** 1 year, 1 month ago

C. Create one AWS SAM configuration file that has default parameters. Perform updates to the testing and staging environments by using the --parameter-overrides flag in the AWS SAM CLI and the parameters that the updates will override.

upvoted 2 times

🗨️ 👤 **NinjaCloud** 1 year, 3 months ago

Correct Answer: C,

You can create a single AWS SAM configuration file with default parameters and then use the --parameter-overrides flag with the AWS SAM CLI to specify parameters that override the defaults for each testing and staging environment. This approach keeps the AWS SAM template file (the infrastructure-as-code) consistent and minimizes duplication. It's a clean and simple way to manage multiple environments without having to create separate templates or custom scripts.

upvoted 8 times

🗨️ 👤 **Rameez1** 1 year, 3 months ago

Selected Answer: C

Here all the options can do the Job but option C does it with least effort.

upvoted 2 times

🗨️ 👤 **PrakashM14** 1 year, 3 months ago

Selected Answer: C

Options A and B introduce additional complexities such as configuration files in TOML format or writing custom shell scripts. These might require more effort and maintenance.

Option D involves adding additional parameters to the existing AWS SAM template, which can work but may lead to a more complex and less maintainable template as the number of environments grows.

Therefore, option C is a straightforward and efficient solution for deploying to multiple environments with AWS SAM.

upvoted 2 times

🗨️ 👤 **dilleman** 1 year, 3 months ago

Selected Answer: A

A is correct

upvoted 4 times

🗨️ 👤 **Digo30sp** 1 year, 3 months ago

Selected Answer: D

The correct answer is (D).

Using the existing AWS SAM template is the option that requires the LEAST development effort. To configure deployments across multiple environments, you can add additional parameters to your AWS SAM template to configure specific attributes for the serverless function and database table resources that are in each environment.

upvoted 3 times

A developer is working on an application that processes operating data from IoT devices. Each IoT device uploads a data file once every hour to an Amazon S3 bucket. The developer wants to immediately process each data file when the data file is uploaded to Amazon S3.

The developer will use an AWS Lambda function to process the data files from Amazon S3. The Lambda function is configured with the S3 bucket information where the files are uploaded. The developer wants to configure the Lambda function to immediately invoke after each data file is uploaded.

Which solution will meet these requirements?


- A. Add an asynchronous invocation to the Lambda function. Select the S3 bucket as the source.
- B. Add an Amazon EventBridge event to the Lambda function. Select the S3 bucket as the source.
- C. Add a trigger to the Lambda function. Select the S3 bucket as the source.
- D. Add a layer to the Lambda function. Select the S3 bucket as the source.

Correct Answer: C

Community vote distribution

C (80%)

B (20%)

 **Digo30sp** Highly Voted 1 year, 3 months ago

Selected Answer: C

The correct answer is (C).

Adding a trigger to your Lambda function is the solution that will meet these requirements. A trigger is an event that can invoke a Lambda function. In the case of this issue, the trigger must be an Amazon S3 event that fires when a new file is uploaded to the bucket.

upvoted 6 times

 **Saurabh04** Most Recent 5 months, 2 weeks ago

Selected Answer: B

Option B is quicker than Option C, because S3 bucket trigger does not guarantee immediate invocation. It relies on event notification from S3.

upvoted 1 times

 **BrainFried** 6 months, 1 week ago

Selected Answer: B


You cannot add a Trigger directly to Lambda. If you want to choose C, then the answer should state: "Add a trigger to S3, select Lambda as the destination"

Since C states "Add trigger to Lambda" (which isn't possible), I will select answer B.

See <https://docs.aws.amazon.com/lambda/latest/dg/lambda-services.html>

"The trigger is actually stored and managed by the service that generates the events, not by Lambda."


upvoted 1 times

 **Anandesh** 6 months, 2 weeks ago

Selected Answer: C

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

upvoted 2 times

 **65703c1** 8 months, 1 week ago

Selected Answer: C

C is the correct answer.

upvoted 1 times

 **1dfed2b** 9 months ago

Selected Answer: B

sure that B, give me a link why everyone want C.

upvoted 1 times

🗨️ 👤 **tsangckl** 7 months, 3 weeks ago
<https://docs.aws.amazon.com/lambda/latest/dg/with-s3-example.html>
upvoted 1 times

🗨️ 👤 **BrainFried** 6 months, 1 week ago
From what I've read, you do not add a trigger to a lambda, you add it else-where (in this case, you add the trigger to S3). The answer says "Add a trigger to Lambda" - this isn't possible!

Read: <https://docs.aws.amazon.com/lambda/latest/dg/lambda-services.html>
"The trigger is actually stored and managed by the service that generates the events, not by Lambda."

The answer should be B then, since EventBridge can monitor S3 bucket and invoke Lambda with the new data.
upvoted 1 times

🗨️ 👤 **SerialiDr** 1 year ago

Selected Answer: C

To meet the requirement of processing data files immediately after they are uploaded to an Amazon S3 bucket, the best solution is to add a trigger to the AWS Lambda function with the S3 bucket as the source. This will configure the Lambda function to be automatically invoked when a new file is uploaded to the specified S3 bucket.

upvoted 3 times

🗨️ 👤 **Certified101** 1 year, 1 month ago

Selected Answer: C

C using S3 Events, no need for EventBridge here.

upvoted 1 times

🗨️ 👤 **LR2023** 1 year, 1 month ago

Selected Answer: B

EventBridge can be employed to collect real-time data streams from various sources like IoT devices, mobile apps, or web applications. Lambda functions can then process this data to perform analytics, generate alerts, or update dashboards.

upvoted 1 times

🗨️ 👤 **LR2023** 1 year, 1 month ago

You can use Amazon EventBridge to monitor an S3 bucket for new image uploads. When a new image is detected, EventBridge triggers a Lambda function that processes the image, applies filters, and generates thumbnails, all without manual intervention

upvoted 1 times

🗨️ 👤 **dilleman** 1 year, 3 months ago

Selected Answer: C

C is correct

upvoted 3 times

🗨️ 👤 **fordiscussiontwo** 1 year, 3 months ago

CCCCCCCCCCCC

upvoted 4 times

A developer is setting up infrastructure by using AWS CloudFormation. If an error occurs when the resources described in the CloudFormation template are provisioned, successfully provisioned resources must be preserved. The developer must provision and update the CloudFormation stack by using the AWS CLI.

Which solution will meet these requirements?

- A. Add a `--enable-termination-protection` command line option to the `create-stack` command and the `update-stack` command.
- B. Add a `--disable-rollback` command line option to the `create-stack` command and the `update-stack` command.
- C. Add a `--parameters ParameterKey=PreserveResources,ParameterValue=True` command line option to the `create-stack` command and the `update-stack` command.
- D. Add a `--tags Key=PreserveResources,Value=True` command line option to the `create-stack` command and the `update-stack` command.

Correct Answer: B

Community vote distribution

B (100%)

 **Digo30sp** Highly Voted 1 year, 3 months ago

Selected Answer: B

The correct answer is (B).

The `--disable-rollback` command-line option will prevent CloudFormation from rolling back the stack to the previous state if an error occurs. This will ensure that successfully provisioned resources are preserved.

upvoted 7 times


 **albert_kuo** Most Recent 4 months ago

Selected Answer: B

```
aws cloudformation create-stack \
--stack-name my-app-stack \
--template-body file://my-template.yaml \
--parameters ParameterKey=InstanceType,ParameterValue=t2.micro \
--disable-rollback
```

```
aws cloudformation update-stack \
--stack-name my-app-stack \
--template-body file://my-template-updated.yaml \
--parameters ParameterKey=InstanceType,ParameterValue=t2.medium \
--disable-rollback
```

upvoted 1 times

 **65703c1** 8 months, 1 week ago

Selected Answer: B

B is the correct answer.

upvoted 1 times

 **KarBiswa** 11 months ago

Selected Answer: B

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stack-failure-options.html>

upvoted 1 times

 **joshnort** 12 months ago

Selected Answer: B

It should look like this:

```
aws cloudformation create-stack --stack-name myteststack --template-body file://DOC-EXAMPLE-BUCKET.json --disable-rollback
```

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stack-failure-options.html#stack-failure-options-cli>
upvoted 3 times

🗨️ 👤 **kaes** 1 year, 2 months ago

Selected Answer: B

"Specify the disable-rollback option or on-failure DO_NOTHING enumeration during a create-stack operation"

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stack-failure-options.html>

upvoted 3 times

🗨️ 👤 **dilleman** 1 year, 3 months ago

Selected Answer: B

B is correct

upvoted 3 times

🗨️ 👤 **kashtelyan** 1 year, 3 months ago

Selected Answer: B

<https://www.cloudhensive.com/blog-posts/cloudformation-disable-rollback/>

upvoted 4 times

🗨️ 👤 **fordiscussiontwo** 1 year, 3 months ago

BBBBBBBBBBBBBBBBBB

upvoted 3 times

A developer is building a serverless application that connects to an Amazon Aurora PostgreSQL database. The serverless application consists of hundreds of AWS Lambda functions. During every Lambda function scale out, a new database connection is made that increases database resource consumption.

The developer needs to decrease the number of connections made to the database. The solution must not impact the scalability of the Lambda functions.

Which solution will meet these requirements?

- A. Configure provisioned concurrency for each Lambda function by setting the ProvisionedConcurrentExecutions parameter to 10.
- B. Enable cluster cache management for Aurora PostgreSQL. Change the connection string of each Lambda function to point to cluster cache management.
- C. Use Amazon RDS Proxy to create a connection pool to manage the database connections. Change the connection string of each Lambda function to reference the proxy.
- D. Configure reserved concurrency for each Lambda function by setting the ReservedConcurrentExecutions parameter to 10.

Correct Answer: C

Community vote distribution

C (100%)

 **dilleman** Highly Voted 9 months, 3 weeks ago

Selected Answer: C

C: Amazon RDS Proxy is designed to improve application scalability and resilience by pooling and reusing database connections. This can significantly reduce the number of connections each Lambda function has to establish
upvoted 9 times

 **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: C

C is the correct answer.
upvoted 1 times

 **Digo30sp** 10 months ago

Selected Answer: C

The correct answer is (C).

Amazon RDS Proxy is a solution that allows you to create a connection pool to manage database connections. This can help reduce the number of connections made to the database.

upvoted 4 times

 **fordiscussiontwo** 10 months ago

CCCCCCCCCCCCCCC

upvoted 4 times

A developer is preparing to begin development of a new version of an application. The previous version of the application is deployed in a production environment. The developer needs to deploy fixes and updates to the current version during the development of the new version of the application. The code for the new version of the application is stored in AWS CodeCommit.

Which solution will meet these requirements?

- A. From the main branch, create a feature branch for production bug fixes. Create a second feature branch from the main branch for development of the new version.
- B. Create a Git tag of the code that is currently deployed in production. Create a Git tag for the development of the new version. Push the two tags to the CodeCommit repository.
- C. From the main branch, create a branch of the code that is currently deployed in production. Apply an IAM policy that ensures no other users can push or merge to the branch.
- D. Create a new CodeCommit repository for development of the new version of the application. Create a Git tag for the development of the new version.

Correct Answer: A

Community vote distribution

A (100%)

65703c1 2 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 2 times

SerialIDr 6 months, 1 week ago

Selected Answer: A

The best option here is:

A. From the main branch, create a feature branch for production bug fixes. Create a second feature branch from the main branch for development of the new version.

Here's why this solution is the most suitable:

Separation of Concerns: Creating separate branches for bug fixes and new feature development ensures that changes made for the current production version and the new version do not interfere with each other. This separation is crucial to avoid introducing new bugs into the production version from the development version.

Continuous Integration and Delivery (CI/CD): This approach supports CI/CD practices. Bug fixes can be developed, tested, and merged into the main branch and deployed without impacting the ongoing development of the new version.

upvoted 2 times

dilleman 9 months, 3 weeks ago

Selected Answer: A

A is a common code version control strategy

upvoted 3 times

Digo30sp 10 months ago

Selected Answer: A

A resposta correta é (A).

Criar uma ramificação de recursos para correções de bugs de produção e uma segunda ramificação de recursos para desenvolvimento da nova versão é a solução que atenderá a esses requisitos.

A primeira ramificação de recursos pode ser usada para corrigir bugs ou implementar atualizações para a versão atual do aplicativo. A segunda ramificação de recursos pode ser usada para desenvolver a nova versão do aplicativo.

upvoted 2 times

  **fordiscussiontwo** 10 months ago

AAAAAAAAAAAAAA

upvoted 3 times

A developer is creating an AWS CloudFormation stack. The stack contains IAM resources with custom names. When the developer tries to deploy the stack, they receive an `InsufficientCapabilities` error.


What should the developer do to resolve this issue?

- A. Specify the `CAPABILITY_AUTO_EXPAND` capability in the CloudFormation stack.
- B. Use an administrators role to deploy IAM resources with CloudFormation.
- C. Specify the `CAPABILITY_IAM` capability in the CloudFormation stack.
- D. Specify the `CAPABILITY_NAMED_IAM` capability in the CloudFormation stack.

Correct Answer: D

Community vote distribution

D (100%)

 **Digo30sp** Highly Voted 1 year, 3 months ago

Selected Answer: D

The correct answer is (D).

To deploy IAM resources with custom names, you must specify the `CAPABILITY_NAMED_IAM` resource in the CloudFormation stack.

The `CAPABILITY_IAM` resource allows CloudFormation to create and modify IAM resources. The `CAPABILITY_NAMED_IAM` resource allows CloudFormation to create IAM resources with custom names.

To resolve the issue, the developer must specify the `CAPABILITY_NAMED_IAM` resource in the CloudFormation stack.

upvoted 9 times

 **Learning4life** Highly Voted 1 year, 3 months ago

D.

If you have IAM resources with custom names, you must specify `CAPABILITY_NAMED_IAM`. See more details in this link

https://docs.aws.amazon.com/AWSCloudFormation/latest/APIReference/API_CreateStack.html

upvoted 5 times

 **albert_kuo** Most Recent 4 months ago

Selected Answer: D


```
aws cloudformation create-stack \
```

```
--stack-name my-iam-stack \
```

```
--template-body file://my-template.yaml \
```

```
--capabilities CAPABILITY_NAMED_IAM
```

upvoted 1 times

 **65703c1** 8 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

 **SerialIDr** 11 months ago

Selected Answer: D

This capability is required when you are deploying IAM resources with custom names using CloudFormation, as it acknowledges that you're creating IAM resources that might affect permissions in your AWS environment.

upvoted 2 times

 **dilleman** 1 year, 3 months ago

Selected Answer: D

D is correct

upvoted 3 times

🗨️ 👤 **Patel_ajay745** 1 year, 3 months ago

CCC

ccccccc

upvoted 1 times

🗨️ 👤 **fordiscussiontwo** 1 year, 3 months ago

DDDDDDDDDD

upvoted 3 times

A company uses Amazon API Gateway to expose a set of APIs to customers. The APIs have caching enabled in API Gateway. Customers need a way to invalidate the cache for each API when they test the API.

What should a developer do to give customers the ability to invalidate the API cache?

- A. Ask the customers to use AWS credentials to call the InvalidateCache API operation.
- B. Attach an InvalidateCache policy to the IAM execution role that the customers use to invoke the API. Ask the customers to send a request that contains the Cache-Control:max-age=0 HTTP header when they make an API call.
- C. Ask the customers to use the AWS SDK API Gateway class to invoke the InvalidateCache API operation.
- D. Attach an InvalidateCache policy to the IAM execution role that the customers use to invoke the API. Ask the customers to add the INVALIDATE_CACHE query string parameter when they make an API call.

Correct Answer: B

Community vote distribution

B (100%)

 **Digo30sp** Highly Voted 10 months ago

Selected Answer: B

B) <https://www.examttopics.com/discussions/amazon/view/4166-exam-aws-certified-developer-associate-topic-1-question-69/>
upvoted 8 times

 **65703c1** Most Recent 2 months, 1 week ago

Selected Answer: B

B is the correct answer.
upvoted 1 times

 **KarBiswa** 5 months ago

Selected Answer: B

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-caching.html#:~:text=A%20client%20of,the%20integration%20endpoint.>
upvoted 2 times

 **KillThemWithKindness** 5 months, 1 week ago

Selected Answer: B

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-caching.html>
upvoted 3 times

 **Mimi666** 8 months ago

invalidate an API Gateway cache entry

A client of your API can invalidate an existing cache entry and reload it from the integration endpoint for individual requests. The client must send a request that contains the Cache-Control: max-age=0 header. The client receives the response directly from the integration endpoint instead of the cache, provided that the client is authorized to do so. This replaces the existing cache entry with the new response, which is fetched from the integration endpoint. <https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-caching.html>
upvoted 2 times

 **dezoito** 9 months, 2 weeks ago

Seems to be B but policies/roles have nothing to do with cache
upvoted 1 times

 **Patel_ajay745** 9 months, 4 weeks ago

it is DDDDD
upvoted 1 times

 **fordiscussiontwo** 9 months, 3 weeks ago

why? because chatGPDUMP said that? all your answers are wrong.
upvoted 8 times

 **fordiscussiontwo** 10 months ago

BBBBBBBBBBBBBB

upvoted 4 times

A developer is creating an AWS Lambda function that will generate and export a file. The function requires 100 MB of temporary storage for temporary files while running. These files will not be needed after the function is complete.

How can the developer MOST efficiently handle the temporary files?

- A. Store the files in Amazon Elastic Block Store (Amazon EBS) and delete the files at the end of the Lambda function.
- B. Copy the files to Amazon Elastic File System (Amazon EFS) and delete the files at the end of the Lambda function.
- C. Store the files in the /tmp directory and delete the files at the end of the Lambda function.
- D. Copy the files to an Amazon S3 bucket with a lifecycle policy to delete the files.

Correct Answer: C

Community vote distribution

C (100%)

🗉 **tsangckl** 7 months, 2 weeks ago

This appear at 17 Jun exam
upvoted 1 times

🗉 **65703c1** 8 months, 1 week ago

Selected Answer: C

C is the correct answer.
upvoted 2 times

🗉 **RPRAMSUBU** 12 months ago

C is the correct answer
upvoted 2 times

🗉 **SerialIDr** 1 year ago

Selected Answer: C

This is the most efficient and straightforward option. AWS Lambda provides a /tmp directory in its execution environment with a storage limit of 512 MB. This space can be used for temporary storage during the function execution. Since the requirement is 100 MB, it falls well within the limits of the /tmp directory.
upvoted 3 times

🗉 **JohnPI** 1 year ago

Selected Answer: C

C is the correct answer
upvoted 1 times

🗉 **sankhagg** 1 year, 1 month ago

Starting March 2022, Lambda now supports increasing /tmp directory's maximum size limit up to 10,240MB. More information available.
<https://aws.amazon.com/blogs/aws/aws-lambda-now-supports-up-to-10-gb-ephemeral-storage/>
upvoted 4 times

🗉 **Claire_KMT** 1 year, 3 months ago

C. Store the files in the /tmp directory and delete the files at the end of the Lambda function.
The /tmp directory is a dedicated temporary storage location provided by AWS Lambda for storing temporary files during the execution of the function.

It's cost-effective and efficient because it doesn't involve additional AWS services or storage costs.

AWS Lambda automatically manages the /tmp directory for you, including clearing its contents after the function execution is complete. You don't need to explicitly delete the files; Lambda takes care of it.

upvoted 3 times

🗉 **LemonGremlin** 1 year, 3 months ago

Selected Answer: C

Option C is the best choice for efficient handling of temporary files within an AWS Lambda function.
upvoted 2 times

A company uses Amazon DynamoDB as a data store for its order management system. The company frontend application stores orders in a DynamoDB table. The DynamoDB table is configured to send change events to a DynamoDB stream. The company uses an AWS Lambda function to log and process the incoming orders based on data from the DynamoDB stream.

An operational review reveals that the order quantity of incoming orders is sometimes set to 0. A developer needs to create a dashboard that will show how many unique customers this problem affects each day.

What should the developer do to implement the dashboard?

- A. Grant the Lambda function's execution role permissions to upload logs to Amazon CloudWatch Logs. Implement a CloudWatch Logs Insights query that selects the number of unique customers for orders with order quantity equal to 0 and groups the results in 1-day periods. Add the CloudWatch Logs Insights query to a CloudWatch dashboard.
- B. Use Amazon Athena to query AWS CloudTrail API logs for API calls. Implement an Athena query that selects the number of unique customers for orders with order quantity equal to 0 and groups the results in 1-day periods. Add the Athena query to an Amazon CloudWatch dashboard.
- C. Configure the Lambda function to send events to Amazon EventBridge. Create an EventBridge rule that groups the number of unique customers for orders with order quantity equal to 0 in 1-day periods. Add a CloudWatch dashboard as the target of the rule.
- D. Turn on custom Amazon CloudWatch metrics for the DynamoDB stream of the DynamoDB table. Create a CloudWatch alarm that groups the number of unique customers for orders with order quantity equal to 0 in 1-day periods. Add the CloudWatch alarm to a CloudWatch dashboard.

Correct Answer: A

Community vote distribution

A (80%)

D (20%)

☰ **Claire_KMT** Highly Voted 1 year, 3 months ago

A. Grant the Lambda function's execution role permissions to upload logs to Amazon CloudWatch Logs. Implement a CloudWatch Logs Insights query that selects the number of unique customers for orders with order quantity equal to 0 and groups the results in 1-day periods. Add the CloudWatch Logs Insights query to a CloudWatch dashboard.

Here's why this option is the best choice:

CloudWatch Logs Insights is designed for querying and analyzing log data, making it well-suited for this task.

By configuring the Lambda function's execution role to upload logs to CloudWatch Logs, you ensure that the log data is available for analysis.

You can use a CloudWatch Logs Insights query to identify unique customers for orders with a quantity of 0 and group the results by day, providing the desired daily count of affected customers.

The results of the query can be added to a CloudWatch dashboard, making it easily accessible for monitoring.

upvoted 5 times

☰ **konieczny69** Highly Voted 12 months ago

Selected Answer: A

D is invalid. There are no such custom metrics:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/metrics-dimensions.html>

A is the right choice

upvoted 5 times

☰ **Saudis** Most Recent 3 months ago

Selected Answer: A

I am with A

upvoted 1 times

🗨️ 👤 **65703c1** 8 months, 1 week ago

Selected Answer: A

A is the correct answer.

upvoted 1 times

🗨️ 👤 **SerialiDr** 1 year ago

Selected Answer: A

This approach involves enhancing the existing Lambda function to log relevant information about orders (especially those with order quantity 0) to CloudWatch Logs. The developer can then use CloudWatch Logs Insights to query these logs for unique customer counts and visualize this data on a CloudWatch dashboard. This solution is feasible and effective, as it leverages the existing Lambda function and CloudWatch capabilities.

upvoted 2 times

🗨️ 👤 **Certified101** 1 year, 1 month ago

Selected Answer: A

<https://www.examttopics.com/discussions/amazon/view/96212-exam-aws-certified-developer-associate-topic-1-question-402/>

upvoted 2 times

🗨️ 👤 **TanTran04** 1 year, 1 month ago

Selected Answer: D

Choose D

Option A is more suitable for log analysis, but in this case, the issue is related to DynamoDB data, and CloudWatch Logs may not be the most efficient way to track it.

upvoted 1 times

🗨️ 👤 **chris_777** 1 year, 2 months ago

Selected Answer: A

I choose A

upvoted 2 times

🗨️ 👤 **bhanupriya07** 1 year, 2 months ago

<https://www.examttopics.com/discussions/amazon/view/96212-exam-aws-certified-developer-associate-topic-1-question-402/>

upvoted 3 times

🗨️ 👤 **PrakashM14** 1 year, 3 months ago

Selected Answer: D

Option A suggests using CloudWatch Logs Insights, which is typically used for analyzing log data. However, in this scenario, the issue is related to metrics (order quantity), and using CloudWatch Metrics and Alarms is a more suitable approach.

I'd go with option D. It seems like a more direct and efficient approach. By using custom CloudWatch metrics for the DynamoDB stream, you can specifically track the relevant data without the need for additional CloudWatch Logs Insights queries. The alarm will then allow you to easily visualize and monitor the number of unique customers affected by the issue each day on the CloudWatch dashboard.

upvoted 2 times

A developer needs to troubleshoot an AWS Lambda function in a development environment. The Lambda function is configured in VPC mode and needs to connect to an existing Amazon RDS for SQL Server DB instance. The DB instance is deployed in a private subnet and accepts connections by using port 1433.


When the developer tests the function, the function reports an error when it tries to connect to the database.

Which combination of steps should the developer take to diagnose this issue? (Choose two.)

- A. Check that the function's security group has outbound access on port 1433 to the DB instance's security group. Check that the DB instance's security group has inbound access on port 1433 from the function's security group.
- B. Check that the function's security group has inbound access on port 1433 from the DB instance's security group. Check that the DB instance's security group has outbound access on port 1433 to the function's security group.
- C. Check that the VPC is set up for a NAT gateway. Check that the DB instance has the public access option turned on.
- D. Check that the function's execution role permissions include `rds:DescribeDBInstances`, `rds:ModifyDBInstance`, and `rds:DescribeDBSecurityGroups` for the DB instance.
- E. Check that the function's execution role permissions include `ec2:CreateNetworkInterface`, `ec2:DescribeNetworkInterfaces`, and `ec2>DeleteNetworkInterface`.

Correct Answer: AE

Community vote distribution



kaes Highly Voted 8 months, 1 week ago

Selected Answer: AE

- A: The function needs outbound access to DB and the DB needs to allow inbound access from the function
- E: The function needs `AWSLambdaVPCLambdaAccessExecutionRole` role to work correctly in the VPC (<https://docs.aws.amazon.com/lambda/latest/dg/configuration-vpc.html#vpc-permissions>)

D is incorrect as the function's execution role does not need to make any of those DB actions: Describe Modify and DescribeDB security groups!
upvoted 10 times

mitch151 Highly Voted 9 months, 1 week ago

I believe it's A and D. Unsure on A, but D seems to be confirmed by this link: <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/lambda-rds-connect.html>

upvoted 8 times

65703c1 Most Recent 2 months, 1 week ago

Selected Answer: AE

AE is the correct answer.
upvoted 1 times

SerialIDr 5 months ago

Selected Answer: AE

A. Check that the function's security group has outbound access on port 1433 to the DB instance's security group. Ensure that the DB instance's security group has inbound access on port 1433 from the function's security group. This setup allows the Lambda function to initiate a connection to the DB instance through the specified port.

E. Check that the function's execution role permissions include `ec2:CreateNetworkInterface`, `ec2:DescribeNetworkInterfaces`, and `ec2>DeleteNetworkInterface`. These permissions are necessary for the Lambda function to create, manage, and clean up the network interfaces that allow it to connect to resources within a VPC, including the RDS instance.

upvoted 2 times

KarBiswa 5 months ago

Selected Answer: AE

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-vpc.html#vpc-permissions>

upvoted 1 times

🗨️ **koniczny69** 6 months ago

Selected Answer: AE

AE

This is a network issue, not a governance issue, hence D is invalid.

Between A and B its an obvious choice.

C is invalid - DB is in a private subnet

upvoted 1 times

🗨️ **_YaWeb** 6 months, 1 week ago

ChatGPT goes with A and D

upvoted 1 times

🗨️ **Snape** 6 months, 3 weeks ago

Selected Answer: AB

inbound and outbound connection between Lambda and the RDS should be set properly.

upvoted 1 times

🗨️ **rrshah83** 7 months ago

Selected Answer: AE

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-vpc.html#vpc-permissions>

upvoted 2 times

🗨️ **Certified101** 7 months, 2 weeks ago

Selected Answer: AE

Agree with Kaes

- A: The function needs outbound access to DB and the DB needs to allow inbound access from the function

- E: The function needs AWSLambdaVPCAccessExecutionRole role to work correctly in the VPC

(<https://docs.aws.amazon.com/lambda/latest/dg/configuration-vpc.html#vpc-permissions>)

D is incorrect as the function's execution role does not need to make any of those DB actions: Describe Modify and DescribeDB security groups!

upvoted 2 times

🗨️ **joshnort** 6 months ago

This is excellent. Thanks for the link. Makes it very clear.

upvoted 1 times

🗨️ **TanTran04** 7 months, 2 weeks ago

Selected Answer: AD

We need connection between lambda and RDS, not to VPC. So, option E is unsuitable. We can choose the related remain option like D

About option A, it's already correct.

upvoted 2 times

🗨️ **Jing2023** 9 months, 1 week ago

Selected Answer: AD

A and D

upvoted 5 times

🗨️ **Claire_KMT** 9 months, 1 week ago

A and B

upvoted 1 times

A developer needs to launch a new Amazon EC2 instance by using the AWS CLI.

Which AWS CLI command should the developer use to meet this requirement?

- A. aws ec2 bundle-instance
- B. aws ec2 start-instances
- C. aws ec2 confirm-product-instance
- D. aws ec2 run-instances

Correct Answer: D

Community vote distribution

D (100%)

☰ **Claire_KMT** Highly Voted 1 year, 3 months ago

D. aws ec2 run-instances

So, to create a new EC2 instance using the AWS CLI, you would typically use the aws ec2 run-instances command, providing the necessary parameters such as the AMI ID, instance type, security groups, and key pair, among others.

upvoted 8 times

☰ **chris_777** Highly Voted 1 year, 2 months ago

Selected Answer: D

D. aws ec2 run-instances

Note: B aws ec2 start-instances is used to "start an instance that you've previously stopped"

upvoted 7 times

☰ **albert_kuo** Most Recent 4 months ago

Selected Answer: D

```
aws ec2 run-instances \  
--image-id ami-xxxxxxx \  
--count 1 \  
--instance-type t2.micro \  
--key-name MyKeyPair \  
--security-group-ids sg-xxxxxxx \  
--subnet-id subnet-xxxxxxx  
upvoted 1 times
```

☰ **65703c1** 8 months, 1 week ago

Selected Answer: D

D is the correct answer.

upvoted 1 times