



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



CERTIFICATION TEST

- CertificationTest.net - Cheap & Quality Resources With Best Support

A company has deployed an e-commerce web application in a new AWS account. An Amazon RDS for MySQL Multi-AZ DB instance is part of this deployment with a database-1.xxxxxxxxxx.us-east-1.rds.amazonaws.com endpoint listening on port 3306. The company's Database Specialist is able to log in to MySQL and run queries from the bastion host using these details.

When users try to utilize the application hosted in the AWS account, they are presented with a generic error message. The application servers are logging a `could not connect to server: Connection times out` error message to Amazon CloudWatch Logs.

What is the cause of this error?

- A. The user name and password the application is using are incorrect.
- B. The security group assigned to the application servers does not have the necessary rules to allow inbound connections from the DB instance.
- C. The security group assigned to the DB instance does not have the necessary rules to allow inbound connections from the application servers.
- D. The user name and password are correct, but the user is not authorized to use the DB instance.

Suggested Answer: C

Reference:

<https://forums.aws.amazon.com/thread.jspa?threadID=129700>

Community vote distribution

C (100%)

aws4myself **Highly Voted** 3 years, 6 months ago

Correct Answer: C

upvoted 7 times

storm0710 **Most Recent** 1 year, 1 month ago

B. Since requirement is of handling random activity peaks

upvoted 1 times

Pranava_GCP 1 year, 9 months ago

Selected Answer: C

C. The security group assigned to the DB instance does not have the necessary rules to allow inbound connections from the application servers.

upvoted 1 times

haison8x 2 years ago

I got this question in the exam:

an aurora cluster has some replica in current region and other regions

the admin accidentally delete primary cluster.

- a. all replica are promoted to standalone
- b. all replica in other regions are promoted
- c. only replica in current region are promoted
- d. dont remember

upvoted 2 times

Kodoma 2 years ago

Do you know what the correct answer was?

upvoted 1 times

chikorita 1 year, 8 months ago

please try to remember option D

upvoted 1 times

novice_expert 3 years ago

Selected Answer: C

DB needs to allow inbound



upvoted 2 times

megadba 3 years, 1 month ago

Selected Answer: C

Correct Answer: C



upvoted 2 times

  **guru_ji** 3 years, 6 months ago

Anyone passed "AWS Certified Database - Specialty" exam recently, could you please share your experience of actual exam??
How many Questions were came from examtopics Q set of 145 in real exam?



Your comments will be highly appreciated.. Thanks..

upvoted 1 times

  **guru_ji** 3 years, 6 months ago

I gave exam today. Only 40 questions came from here out of 65.

upvoted 2 times

  **lihze** 3 years, 6 months ago

Guru - 40 questions from out of 112 or 149? I can able to access only 112. Can you possible to share the dump if you have 149 questions.



upvoted 2 times

  **aws4myself** 3 years, 6 months ago

Basically, it is a connection timeout error, so A, D can be deleted.

B is wrong, because if it has a problem in SG rules, it should affect at outbound not inbound,
Hence C is correct, which is the inbound rule for DB instance SG.



upvoted 1 times

  **guru_ji** 3 years, 6 months ago

Are you planning for exam?

We can share study material, it would be beneficial for both. You can email me on "awsdbguru at gmail"

upvoted 1 times

  **guru_ji** 3 years, 6 months ago



Answer is C

Anyone preparing for DB Speciality and want to do group study with us, comment below with email
upvoted 1 times

  **anandbabu** 3 years, 6 months ago



its C for me

upvoted 2 times

  **guru_ji** 3 years, 6 months ago

C ==>> Correct Answer.

upvoted 1 times

  **Dr_Kiko** 3 years, 6 months ago

It's C. I actually had to configure SG on my MYSQL instance to allow me connect from home PC with a SQL client :)

upvoted 1 times

  **db_interest** 3 years, 7 months ago

A company has an application that uses an Amazon DynamoDB table to store user dat

a. Every morning, a single-threaded process calls the DynamoDB API Scan operation to scan the entire table and generate a critical start-of-day report for management. A successful marketing campaign recently doubled the number of items in the table, and now the process takes too long to run and the report is not generated in time. A database specialist needs to improve the performance of the process. The database specialist notes that, when the process is running, 15% of the table's provisioned read capacity units (RCUs) are being used.

What should the database specialist do?

- A. Enable auto scaling for the DynamoDB table.
- B. Use four threads and parallel DynamoDB API Scan operations.
- C. Double the table's provisioned RCUs.
- D. Set the Limit and Offset parameters before every call to the API.

upvoted 1 times

  **Zhongkai** 3 years, 7 months ago

Looks like B

upvoted 1 times

🗨️ 👤 **db_interest** 3 years, 6 months ago

Yep, B looks right. Since autoscaling would increase the WCU/RCU which is not needed.
upvoted 1 times

🗨️ 👤 **db_interest** 3 years, 7 months ago

A company is running an on-premises application comprised of a web tier, an application tier, and a MySQL database tier. The database is used primarily during business hours with random activity peaks throughout the day. A database specialist needs to improve the availability and reduce the cost of the MySQL database tier as part of the company's migration to AWS.

Which MySQL database option would meet these requirements?

- A. Amazon RDS for MySQL with Multi-AZ
- B. Amazon Aurora Serverless MySQL cluster
- C. Amazon Aurora MySQL cluster
- D. Amazon RDS for MySQL with read replica

upvoted 1 times

🗨️ 👤 **Zhongkai** 3 years, 7 months ago

random activity peaks -> Aurora Serverless -> B

upvoted 2 times

🗨️ 👤 **ch321** 3 years, 5 months ago

Should be B

upvoted 1 times

🗨️ 👤 **agrawalachin** 3 years, 7 months ago

A company is running an on-premises application comprised of a web tier, an application tier, and a MySQL database tier. The database is used primarily during business hours with random activity peaks throughout the day. A database specialist needs to improve the availability and reduce the cost of the MySQL database tier as part of the company's migration to AWS.

Which MySQL database option would meet these requirements?

- A. Amazon RDS for MySQL with Multi-AZ
- B. Amazon Aurora Serverless MySQL cluster
- C. Amazon Aurora MySQL cluster
- D. Amazon RDS for MySQL with read replica

upvoted 1 times

🗨️ 👤 **agrawalachin** 3 years, 7 months ago

A database specialist needs to review and optimize an Amazon DynamoDB table that is experiencing performance issues. A thorough investigation by the database specialist reveals that the partition key is causing hot partitions, so a new partition key is created. The database specialist must effectively apply this new partition key to all existing and new data.

How can this solution be implemented?

- A. Use Amazon EMR to export the data from the current DynamoDB table to Amazon S3. Then use Amazon EMR again to import the data from Amazon S3 into a new DynamoDB table with the new partition key.
- B. Use AWS DMS to copy the data from the current DynamoDB table to Amazon S3. Then import the DynamoDB table to create a new DynamoDB table with the new partition key.
- C. Use the AWS CLI to update the DynamoDB table and modify the partition key.
- D. Use the AWS CLI to back up the DynamoDB table. Then use the restore-table-from-backup command and modify the partition key.

upvoted 1 times

🗨️ 👤 **db_interest** 3 years, 7 months ago

B. Since requirement is of handling random activity peaks

upvoted 1 times

🗨️ 👤 **db_interest** 3 years, 7 months ago

Ignore the above answer. It's for another question :)

upvoted 1 times



🗨️ 👤 **Zhongkai** 3 years, 7 months ago

DDB cannot change Partition key of a table -> C/D wrong.

DMS cannot choose DDB as source -> B wrong.

So A should be the correct one although it appears to need a lot of efforts.

upvoted 4 times

  **LMax** 3 years, 7 months ago

It's C

upvoted 1 times

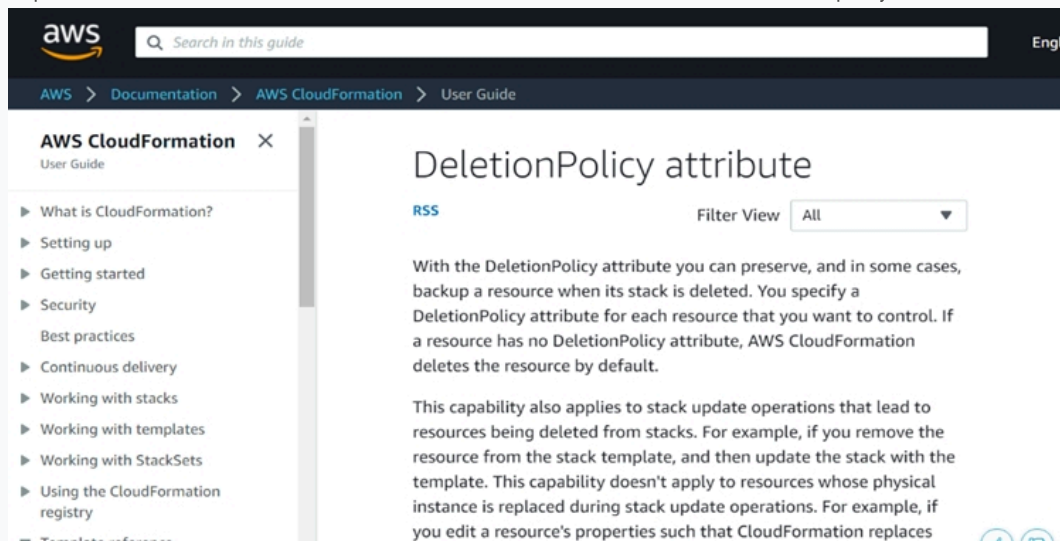
An AWS CloudFormation stack that included an Amazon RDS DB instance was accidentally deleted and recent data was lost. A Database Specialist needs to add RDS settings to the CloudFormation template to reduce the chance of accidental instance data loss in the future. Which settings will meet this requirement? (Choose three.)

- A. Set DeletionProtection to True
- B. Set MultiAZ to True
- C. Set TerminationProtection to True
- D. Set DeleteAutomatedBackups to False
- E. Set DeletionPolicy to Delete
- F. Set DeletionPolicy to Retain

Suggested Answer: ACF

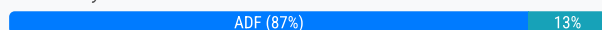
Reference:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html>



<https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-accidental-updates/>

Community vote distribution



🗳️ **jove** Highly Voted 3 years, 5 months ago

Selected Answer: ADF

A - <https://aws.amazon.com/about-aws/whats-new/2018/09/amazon-rds-now-provides-database-deletion-protection/>

D - https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithAutomatedBackups.html

F - <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html>

upvoted 15 times

🗳️ **Sathish_dbs** Most Recent 1 year, 7 months ago

got this Q in the exam

upvoted 1 times

🗳️ **f__16** 1 year, 10 months ago

Selected Answer: ADF

ADF

according to this link <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-rds-dbinstance.html>

upvoted 1 times

🗳️ **sirfans** 2 years, 7 months ago

Selected Answer: ADF

ADF is the right one

upvoted 1 times

🗳️ 👤 **Luke97** 2 years, 8 months ago

TerminationProtection is a delete protection on CloudFormation instead of RDS. So ADF are correct answers.

upvoted 1 times

🗳️ 👤 **sachin** 2 years, 10 months ago

ADF is right . Deletion Policy ;- retain in for any resource attached to RDS .

upvoted 1 times

🗳️ 👤 **Balki** 2 years, 11 months ago

Selected Answer: ADF

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-properties-rds-database-instance.html>

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: ADF

A. Set DeletionProtection to True

x B. Set MultiAZ to True

x C. Set TerminationProtection to True (it is for stack deletion)

D. Set DeleteAutomatedBackups to False

x E. Set DeletionPolicy to Delete

F. Set DeletionPolicy to Retain

upvoted 1 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: ADF

Very frequently asked question used in different question banks.

Set DeletionProtection of the Stack

Prevent Automated Backups from being deleted

Set deletion policy to Retain (default is snapshot)

upvoted 3 times

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

The default of DeletionPolicy may be "delete", which delete the resources

upvoted 1 times

🗳️ 👤 **jeyp12** 3 years, 3 months ago

Selected Answer: CDF

Are we not applying this to stack so should be C,D,F..not to RDS..there is no deletion protection at stack level

upvoted 4 times

🗳️ 👤 **novice_expert** 3 years ago

ADF

<https://aws.amazon.com/about-aws/whats-new/2017/09/aws-cloudformation-provides-stack-termination-protection/#:~:text=AWS%20CloudFormation%20now%20allows%20you,its%20status%2C%20will%20remain%20unchanged.>

upvoted 1 times

🗳️ 👤 **szl0144** 3 years, 4 months ago

the answer is ADF

upvoted 1 times

🗳️ 👤 **victornj** 3 years, 4 months ago

C, D & F

upvoted 1 times

🗳️ 👤 **2025flakyt** 3 years, 5 months ago

ACF is the correct answer.

You need Set TerminationProtection to True for the RDS instance

upvoted 3 times

🗳️ 👤 **jove** 3 years, 5 months ago


RDS doesn't have TerminationProtection .. It has DeletionProtection. So, the correct options are A,D,F

upvoted 3 times

🗳️ 👤 **Pranava_GCP** 1 year, 8 months ago

TerminationProtection setting applies to stack instead of RDS instance.

upvoted 1 times

  **GMartinelli** 3 years, 5 months ago

Selected Answer: ADF

Option A, D, F

RDS doesn't have termination protection, only deletion protection

upvoted 4 times

A Database Specialist is troubleshooting an application connection failure on an Amazon Aurora DB cluster with multiple Aurora Replicas that had been running with no issues for the past 2 months. The connection failure lasted for 5 minutes and corrected itself after that. The Database Specialist reviewed the Amazon RDS events and determined a failover event occurred at that time. The failover process took around 15 seconds to complete. What is the MOST likely cause of the 5-minute connection outage?

- A. After a database crash, Aurora needed to replay the redo log from the last database checkpoint
- B. The client-side application is caching the DNS data and its TTL is set too high
- C. After failover, the Aurora DB cluster needs time to warm up before accepting client connections
- D. There were no active Aurora Replicas in the Aurora DB cluster

Suggested Answer: C

Community vote distribution

B (100%)

🗳️ 👤 **jnassp1** Highly Voted 3 years, 8 months ago

Answer: B

The question indicates failover has already taken place. The TTL is set too high.

upvoted 17 times

🗳️ 👤 **LMax** 3 years, 6 months ago

Agree, B. TTL.

upvoted 3 times

🗳️ 👤 **58a2d17** Most Recent 1 year, 2 months ago

examtopics.com : please explain your answers C.

upvoted 1 times

🗳️ 👤 **Pranava_GCP** 1 year, 8 months ago

Selected Answer: B

B. The client-side application is caching the DNS data and its TTL is set too high

upvoted 1 times

🗳️ 👤 **renfdo** 2 years, 4 months ago

Selected Answer: B

Answer: B

<https://aws.amazon.com/premiumsupport/knowledge-center/aurora-mysql-redirected-endpoint-writer/>

upvoted 1 times

🗳️ 👤 **SteveMartin9** 2 years, 4 months ago

Selected Answer: B

B is the answer

upvoted 1 times

🗳️ 👤 **SachinGoel** 2 years, 4 months ago

Selected Answer: B

B is the answer

upvoted 1 times

🗳️ 👤 **ftrimmer** 2 years, 6 months ago

B is corret

upvoted 1 times

🗳️ 👤 **sirfans** 2 years, 7 months ago

Selected Answer: B

B is the right one.

upvoted 1 times

🗨️ 👤 **niteshdba** 2 years, 8 months ago

B is correct

upvoted 1 times

🗨️ 👤 **novice_expert** 3 years ago

Selected Answer: B

B. The client-side application is caching the DNS data and its TTL is set too high

If your client application is caching the Domain Name Service (DNS) data of your DB instances, set a time-to-live (TTL) value of less than 30 seconds. Because the underlying IP address of a DB instance can change after a failover, caching the DNS data for an extended time can lead to connection failures if your application tries to connect to an IP address that no longer is in service.

upvoted 2 times

🗨️ 👤 **RotterDam** 3 years, 2 months ago

Got this question in my exam. (i cleared it). B is correct

upvoted 2 times

🗨️ 👤 **Zdujgfr567783ff** 2 years, 1 month ago

how would you know that?

in the result there is passed / not passed and a section %%

upvoted 1 times

🗨️ 👤 **kped21** 3 years, 3 months ago

B

If your client application is caching the Domain Name Service (DNS) data of your DB instances, set a time-to-live (TTL) value of less than 30 seconds. Because the underlying IP address of a DB instance can change after a failover, caching the DNS data for an extended time can lead to connection failures if your application tries to connect to an IP address that no longer is in service.

upvoted 2 times

🗨️ 👤 **aws4myself** 3 years, 6 months ago

B is the correct answer, cache needs to be set for the lowest possible time.

upvoted 1 times

🗨️ 👤 **guru_ji** 3 years, 6 months ago

B ==>> Correct Answer

upvoted 1 times

🗨️ 👤 **Dr_Kiko** 3 years, 6 months ago

B; in Stephane's course he talks about Aurora-Postgres settings to keep client caching at minimum. Since there were Read Replicas in place, Aurora-side failover is <1 min

upvoted 1 times

🗨️ 👤 **Dip11** 3 years, 6 months ago

Why "Reveal Solution" answer is mostly wrong ? When i come to discussion i see the answers given by others which look correct/logical. Is "Reveal solution" not reliable ?

upvoted 2 times

🗨️ 👤 **jove** 3 years, 5 months ago

"Reveal solution" is not reliable. Always read the discussion.

upvoted 2 times

🗨️ 👤 **myutran** 3 years, 6 months ago

Ans: B

upvoted 1 times

A company is deploying a solution in Amazon Aurora by migrating from an on-premises system. The IT department has established an AWS Direct Connect link from the company's data center. The company's Database Specialist has selected the option to require SSL/TLS for connectivity to prevent plaintext data from being sent over the network. The migration appears to be working successfully, and the data can be queried from a desktop machine.

Two Data Analysts have been asked to query and validate the data in the new Aurora DB cluster. Both Analysts are unable to connect to Aurora. Their user names and passwords have been verified as valid and the Database Specialist can connect to the DB cluster using their accounts. The Database Specialist also verified that the security group configuration allows network from all corporate IP addresses.

What should the Database Specialist do to correct the Data Analysts' inability to connect?

- A. Restart the DB cluster to apply the SSL change.
- B. Instruct the Data Analysts to download the root certificate and use the SSL certificate on the connection string to connect.
- C. Add explicit mappings between the Data Analysts' IP addresses and the instance in the security group assigned to the DB cluster.
- D. Modify the Data Analysts' local client firewall to allow network traffic to AWS.

Suggested Answer: D

Community vote distribution

B (100%)

🗳️ 👤 **Pranava_GCP** 1 year, 8 months ago

Selected Answer: B

B. Instruct the Data Analysts to download the root certificate and use the SSL certificate on the connection string to connect.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/ssl-certificate-rotation-aurora-postgresql.html>

upvoted 3 times

🗳️ 👤 **SteveMartin9** 2 years, 4 months ago

Author from the Udemy.com practice test says B is the correct answer.

upvoted 4 times

🗳️ 👤 **sirfans** 2 years, 7 months ago

Selected Answer: B

B is the right one

upvoted 1 times

🗳️ 👤 **niteshdba** 2 years, 8 months ago

B is the answer

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: B

B. Instruct the Data Analysts to download the root certificate and use the SSL certificate on the connection string to connect.

To connect using SSL:

- Provide the SSLTrust certificate (can be downloaded from AWS)
- Provide SSL options when connecting to database
- Not using SSL on a DB that enforces SSL would result in error

B - Need root certificate and then need to specify --sql-ca = cert.pem --ssl-mode=verify_identity for example mysql

When the require_secure_transport parameter is set to ON for a DB cluster, a database client can connect to it if it can establish an encrypted connection. Otherwise, an error message similar to the following is returned to the client:

MySQL Error 3159 (HY000): Connections using insecure transport are prohibited while --require_secure_transport=ON.

upvoted 2 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

The answer is (B). what I am beginning to REALLY dislike about some of these questions is the terrible grammar. This question is very similar to another question bank where the Database specialist is able to connect to the Aurora Cluster. The english they are using is "Specialist MAY use their account to log in" - this is artificially vague. in the pressure of time - such english just leaves a bad taste and its clear many of the questions are not written by native english language speakers - which makes it frustrating

upvoted 3 times

🗳️ 👤 **Shunpin** 3 years, 5 months ago

Selected Answer: B

My point: Usually, you have no privileges to modify local firewall policy in a big cooperate. In the question, it has mentioned the connection can be made from desktop that means local firewall rules allow to access RDS.

upvoted 3 times

🗳️ 👤 **aws4myself** 3 years, 6 months ago

its B, SSL issue.

upvoted 1 times

🗳️ 👤 **guru_ji** 3 years, 6 months ago

Correct Answer ==>> B

upvoted 1 times

🗳️ 👤 **Dr_Kiko** 3 years, 6 months ago

B, you need a cert

upvoted 1 times

🗳️ 👤 **LMax** 3 years, 6 months ago

Must be B, SSL issue.

upvoted 3 times

🗳️ 👤 **jnassp1** 3 years, 6 months ago

B

- To connect using SSL:
 - Provide the SSLTrust certificate (can be downloaded from AWS)
 - Provide SSL options when connecting to database
 - Not using SSL on a DB that enforces SSL would result in error
- upvoted 1 times

🗳️ 👤 **jnassp1** 3 years, 7 months ago

D is not right this. The questions is on SSL/TLS encryption in transit -

B - Need root certificate and then need to specify --sql-ca = cert.pem --ssl-mode=verify_identity for example mysql

upvoted 1 times

🗳️ 👤 **myutran** 3 years, 7 months ago

Ans: D

upvoted 1 times

🗳️ 👤 **Exia** 3 years, 7 months ago

When the require_secure_transport parameter is set to ON for a DB cluster, a database client can connect to it if it can establish an encrypted connection. Otherwise, an error message similar to the following is returned to the client:

MySQL Error 3159 (HY000): Connections using insecure transport are prohibited while --require_secure_transport=ON.

upvoted 1 times

🗳️ 👤 **Exia** 3 years, 7 months ago

D. Aurora MySQL DB clusters must be created in an Amazon Virtual Private Cloud (VPC). To control which devices and Amazon EC2 instances can open connections to the endpoint and port of the DB instance for Aurora MySQL DB clusters in a VPC, you use a VPC security group. These endpoint and port connections can be made using Secure Sockets Layer (SSL). In addition, firewall rules at your company can control whether devices running at your company can open connections to a DB instance.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Security.html>

upvoted 2 times

🗳️ 👤 **goodh32** 3 years, 7 months ago

Answer is B

As SSL parameter can be used in string

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/ssl-certificate-rotation-aurora-postgresql.html>
upvoted 3 times

A company is concerned about the cost of a large-scale, transactional application using Amazon DynamoDB that only needs to store data for 2 days before it is deleted. In looking at the tables, a Database Specialist notices that much of the data is months old, and goes back to when the application was first deployed.

What can the Database Specialist do to reduce the overall cost?

- A. Create a new attribute in each table to track the expiration time and create an AWS Glue transformation to delete entries more than 2 days old.
- B. Create a new attribute in each table to track the expiration time and enable DynamoDB Streams on each table.
- C. Create a new attribute in each table to track the expiration time and enable time to live (TTL) on each table.
- D. Create an Amazon CloudWatch Events event to export the data to Amazon S3 daily using AWS Data Pipeline and then truncate the Amazon DynamoDB table.

Suggested Answer: A

Community vote distribution

C (100%)

 **chicagomassageseeker** Highly Voted 3 years, 8 months ago

Answer C. Enable TTL on a new attribute.

upvoted 13 times

 **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: C

C. Create a new attribute in each table to track the expiration time and enable time to live (TTL) on each table.

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


upvoted 2 times

 **cat_of_meerkat** 1 year, 8 months ago

Selected Answer: C

Enable TTL on a new attribute.

upvoted 1 times

 **ken_test1234** 2 years, 2 months ago

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/time-to-live-ttl-before-you-start.html> correct answer C


upvoted 1 times

 **pintu143** 2 years, 3 months ago

Selected Answer: C

Answer C. Enable TTL on a new attribute.

upvoted 1 times

 **renfdo** 2 years, 4 months ago

Selected Answer: C

Answer:C


upvoted 2 times

 **SteveMartin9** 2 years, 4 months ago

Selected Answer: C

Author from the Udemy.com practice test says C is the correct answer.

upvoted 2 times

 **lollyj** 2 years, 5 months ago

Can someone please explain why A was a better solution for cost?

upvoted 1 times

 **ryuhei** 2 years, 11 months ago

Selected Answer: C

Answer:C

upvoted 4 times

🗲️ 👤 **novice_expert** 3 years ago

Selected Answer: C

C. Create a new attribute in each table to track the expiration time and enable time to live (TTL) on each table.

upvoted 3 times

🗲️ 👤 **Anuragdba** 3 years, 6 months ago

C: TTL is used in dynamo DB

upvoted 1 times

🗲️ 👤 **aws4myself** 3 years, 6 months ago

its C, Dynamo DB TTL attribute for auto deletion

upvoted 1 times

🗲️ 👤 **guru_ji** 3 years, 6 months ago

C ==>> Correct Answer.

upvoted 1 times

🗲️ 👤 **Ninjamonkey8812** 3 years, 7 months ago

My opinion:

You can include AWS glue to write a Deletion job to transform the data.

TTL might be bound by the timefactor in the document

upvoted 2 times

🗲️ 👤 **aws4myself** 3 years, 6 months ago

Super, I guess you are an aws product developer/architect. Nice answers, dont mislead people.

upvoted 2 times

🗲️ 👤 **rlnd2000** 2 years, 10 months ago

I don't think your solution helps to minimize the total cost of ownership, the delete queries will scan the tables we will need to pay for that, and we will need to pay for Glue Jobs executions and aws says "...TTL is provided at no extra cost as a means to reduce stored data volumes.."

upvoted 2 times

🗲️ 👤 **Dip11** 3 years, 7 months ago

Yes C is correct

upvoted 1 times

🗲️ 👤 **LMax** 3 years, 7 months ago

For sure Answer C

upvoted 1 times

🗲️ 👤 **myutran** 3 years, 7 months ago

Ans: C

upvoted 1 times

A company has an on-premises system that tracks various database operations that occur over the lifetime of a database, including database shutdown, deletion, creation, and backup.

The company recently moved two databases to Amazon RDS and is looking at a solution that would satisfy these requirements. The data could be used by other systems within the company.

Which solution will meet these requirements with minimal effort?

- A. Create an Amazon CloudWatch Events rule with the operations that need to be tracked on Amazon RDS. Create an AWS Lambda function to act on these rules and write the output to the tracking systems.
- B. Create an AWS Lambda function to trigger on AWS CloudTrail API calls. Filter on specific RDS API calls and write the output to the tracking systems.
- C. Create RDS event subscriptions. Have the tracking systems subscribe to specific RDS event system notifications.
- D. Write RDS logs to Amazon Kinesis Data Firehose. Create an AWS Lambda function to act on these rules and write the output to the tracking systems.

Suggested Answer: C

Community vote distribution


C (100%)

 **chicagomassageseeker** Highly Voted 3 years, 8 months ago

Answer C:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_Events.html

upvoted 8 times

 **ken_test1234** Most Recent 2 years, 2 months ago

Answer C

upvoted 1 times

 **awsguys** 3 years ago

NOT A ==> an on-premises system

upvoted 1 times

 **novice_expert** 3 years ago

Selected Answer: C

Create RDS event subscriptions. Have the tracking systems subscribe to specific RDS event system notifications.


https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_Events.html

upvoted 2 times

 **awsmonster** 3 years, 4 months ago

Shouldn't it be A since the tracking system sits on premise and C does not cater to write the output to the tracking system ?

upvoted 1 times

 **awsmonster** 3 years, 3 months ago

Should be C, after restudying RDS Event

upvoted 1 times

 **GMartinelli** 3 years, 5 months ago

Selected Answer: C


Option C

upvoted 2 times

 **aws4myself** 3 years, 6 months ago

B also can be the answer, but question is with minimal effort, Hence the answer is C

upvoted 2 times

 **guru_ji** 3 years, 7 months ago

Theek answer hai ==>> C

upvoted 1 times

🗨️ 👤 **Dr_Kiko** 3 years, 7 months ago

Answer C

upvoted 1 times

🗨️ 👤 **LMax** 3 years, 7 months ago

Looks like C is the one

upvoted 2 times

🗨️ 👤 **myutran** 3 years, 7 months ago

Ans: C

upvoted 1 times

🗨️ 👤 **Ebi** 3 years, 7 months ago

C is correct

upvoted 2 times

🗨️ 👤 **BillyC** 3 years, 8 months ago

C its correct

upvoted 4 times

A clothing company uses a custom ecommerce application and a PostgreSQL database to sell clothes to thousands of users from multiple countries. The company is migrating its application and database from its on-premises data center to the AWS Cloud. The company has selected Amazon EC2 for the application and Amazon RDS for PostgreSQL for the database. The company requires database passwords to be changed every 60 days. A Database Specialist needs to ensure that the credentials used by the web application to connect to the database are managed securely.

Which approach should the Database Specialist take to securely manage the database credentials?

- A. Store the credentials in a text file in an Amazon S3 bucket. Restrict permissions on the bucket to the IAM role associated with the instance profile only. Modify the application to download the text file and retrieve the credentials on start up. Update the text file every 60 days.
- B. Configure IAM database authentication for the application to connect to the database. Create an IAM user and map it to a separate database user for each ecommerce user. Require users to update their passwords every 60 days.
- C. Store the credentials in AWS Secrets Manager. Restrict permissions on the secret to only the IAM role associated with the instance profile. Modify the application to retrieve the credentials from Secrets Manager on start up. Configure the rotation interval to 60 days.
- D. Store the credentials in an encrypted text file in the application AMI. Use AWS KMS to store the key for decrypting the text file. Modify the application to decrypt the text file and retrieve the credentials on start up. Update the text file and publish a new AMI every 60 days.

Suggested Answer: B

Community vote distribution

C (100%)

 **chicagomassageseeker** Highly Voted 3 years, 8 months ago

Answer C.

upvoted 24 times

 **novice_expert** Highly Voted 3 years ago

Selected Answer: C

Secret Manager

-> rotation 60 days

->Secret access to IAM roles for instance only

-> Apps refer Secret manager to get pwd on startup

upvoted 5 times

 **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: C

C. Store the credentials in AWS Secrets Manager. Restrict permissions on the secret to only the IAM role associated with the instance profile. Modify the application to retrieve the credentials from Secrets Manager on start up. Configure the rotation interval to 60 days

<https://aws.amazon.com/secrets-manager/>

upvoted 2 times

 **kerl** 1 year, 9 months ago

now the answer is B, <https://repost.aws/knowledge-center/users-connect-rds-iam>, "If your application is running on Amazon Elastic Compute Cloud (Amazon EC2), then you can use your EC2 instance profile credentials to access the database. You don't need to store database passwords on your instance."

and "Authentication tokens have a lifespan of 15 minutes, so you don't need to enforce password resets." meet the question criteria.

C no longer best practice.

upvoted 2 times

 **Balki** 2 years, 11 months ago

Selected Answer: C

If people think of B, the only reason they should move away from it is, IAM DB Authentication tokens can be valid only for 15 mins.

Answer is C

upvoted 2 times

🗨️ **amitkhurana** 3 years, 1 month ago

Selected Answer: C

Answer C.

upvoted 2 times

🗨️ **RotterDam** 3 years, 2 months ago

Selected Answer: C

Obviously C. How do the owners allow 80% wrong answers and not correct them?

upvoted 3 times

🗨️ **tugboat** 3 years, 3 months ago

Selected Answer: C

secret manager

upvoted 2 times

🗨️ **Raj12131** 3 years, 4 months ago

B seems to be right choice. <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.IAMDBAuth.html>

upvoted 1 times

🗨️ **Jiang_aws1** 2 years, 7 months ago

X B. Create IAM role then grant to user , not create "IAM user"

upvoted 1 times

🗨️ **GMartinelli** 3 years, 5 months ago

Selected Answer: C

Option C

upvoted 2 times

🗨️ **Anuragdba** 3 years, 6 months ago

C: Why most of answer is wrong ? this is not right way .

upvoted 1 times

🗨️ **Anuragdba** 3 years, 7 months ago

C . Store the credentials in AWS Secrets Manager

upvoted 1 times

🗨️ **aws4myself** 3 years, 7 months ago

C ==> for centralised credentials management with auto rotation

upvoted 2 times

🗨️ **guru_ji** 3 years, 7 months ago

C ==>> Correct Answer.

upvoted 1 times

🗨️ **Dr_Kiko** 3 years, 7 months ago

C

I cannot believe how easy the question is!

upvoted 1 times

🗨️ **Billhardy** 3 years, 7 months ago

Answer C

upvoted 1 times

🗨️ **LMax** 3 years, 7 months ago

Answer C

upvoted 2 times

A financial services company is developing a shared data service that supports different applications from throughout the company. A Database Specialist designed a solution to leverage Amazon ElastiCache for Redis with cluster mode enabled to enhance performance and scalability. The cluster is configured to listen on port 6379.

Which combination of steps should the Database Specialist take to secure the cache data and protect it from unauthorized access? (Choose three.)

- A. Enable in-transit and at-rest encryption on the ElastiCache cluster.
- B. Ensure that Amazon CloudWatch metrics are configured in the ElastiCache cluster.
- C. Ensure the security group for the ElastiCache cluster allows all inbound traffic from itself and inbound traffic on TCP port 6379 from trusted clients only.
- D. Create an IAM policy to allow the application service roles to access all ElastiCache API actions.
- E. Ensure the security group for the ElastiCache clients authorize inbound TCP port 6379 and port 22 traffic from the trusted ElastiCache cluster's security group.
- F. Ensure the cluster is created with the auth-token parameter and that the parameter is used in all subsequent commands.

Suggested Answer: ABE

Reference:

<https://aws.amazon.com/getting-started/tutorials/setting-up-a-redis-cluster-with-amazon-elasticache/>

Setting up a Redis Cluster for scalability and high availability
Amazon ElastiCache for Redis cluster mode setup, configuration, security, and provisioning

In this tutorial, you will learn how to create and configure a Redis Cluster with ElastiCache for Redis. With cluster mode enabled, your Redis Cluster gains enhanced scalability and high availability. You can start small and easily scale your Redis data as your application grows, and by setting up replicas in different availability zones you can also increase your read capacity. By following this tutorial you will also learn some best practices for Redis Cluster deployments.

There's a very rich ecosystem of Redis clients with support for cluster mode features. That makes working with sharded data extremely easy and reliable as clients can keep track of the location of hashslots in a local cache, thus improving the performance of the system as a whole.

The ElastiCache node created in this tutorial is free tier eligible.

About this Tutorial	
Time	10 - 20 minutes
Cost	Free Tier Eligible
Use Case	Scaling, High Availability, Real-time application
Products	AWS ElastiCache for Redis, AWS Free Tier
Audience	Developers

Community vote distribution
ACF (100%)

Pranava_GCP 1 year, 8 months ago

Selected Answer: ACF

ACF are correct

A refer to

<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/encryption.html>

F refers to

<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/auth.html>

upvoted 2 times

mraronsimon 1 year, 10 months ago

Selected Answer: ACF

ACF is correct because those are preventive protection

B - not preventive

D - what about non-IAM-based connections? :)

E - client-server architecture -> always the client initiates the connection! The cluster needs SG to prevent connections from unexpected sources and on unexpected ports, instead of clients...

upvoted 1 times

🗨️ 👤 **ken_test1234** 2 years, 2 months ago

ACF is the correct answer

upvoted 1 times

🗨️ 👤 **SteveMartin9** 2 years, 4 months ago

Selected Answer: ACF

Author from the Udemy.com practice test says ACF is the correct answer.

upvoted 3 times

🗨️ 👤 **SachinGoel** 2 years, 4 months ago

Selected Answer: ACF

ACF is right choice

upvoted 1 times

🗨️ 👤 **sju** 2 years, 5 months ago

Why A, if data is encrypted, it will remain confidential but open for manipulation as you can delete it. Encryption can give confidentiality but can't guarantee integrity.

upvoted 1 times

🗨️ 👤 **Sab** 2 years, 5 months ago

Elasticache for Redis 7.0 now support IAM authentication through users and roles.

upvoted 1 times

🗨️ 👤 **novice_expert** 3 years ago

Selected Answer: ACF

A. Enable in-transit and at-rest encryption on the ElastiCache cluster.

x B. why CloudWatch ?

C. Ensure the security group for the ElastiCache cluster allows all inbound traffic from itself and inbound traffic on TCP port 6379 from trusted clients only.

x why all API? D. Create an IAM policy to allow the application service roles to access all ElastiCache API actions.

x why 22? E. Ensure the security group for the ElastiCache clients authorize inbound TCP port 6379 and port 22 traffic from the trusted ElastiCache cluster's security group.

F. Ensure the cluster is created with the auth-token parameter and that the parameter is used in all subsequent commands.

upvoted 4 times

🗨️ 👤 **soyyodario** 3 years, 4 months ago

Selected Answer: ACF

ACF are the correct

E Why do you need port 22?

upvoted 2 times

🗨️ 👤 **2025flakyt** 3 years, 5 months ago

ADF are the correct options

upvoted 1 times

🗨️ 👤 **2025flakyt** 3 years, 5 months ago

Ensure the security group for the ElastiCache cluster allows all inbound traffic from itself is only needed when you launched your ElastiCache instance in EC2 Classic. so C is not a valid option

upvoted 1 times

🗨️ 👤 **jove** 3 years, 5 months ago

These questions are not for up to date versions. When this question was added the most likely the EC2-Classical was still very much available.

My choice is ACF

upvoted 1 times

🗨️ 👤 **2025flakyt** 3 years, 5 months ago

The following is needed to protect ElastiCache

Use multi-factor authentication (MFA) with each account.

Use SSL/TLS to communicate with AWS resources.

Set up API and user activity logging with AWS CloudTrail.

Use AWS encryption solutions, along with all default security controls within AWS services.

Use advanced managed security services such as Amazon Macie, which assists in discovering and securing personal data that is stored in Amazon S3.

upvoted 2 times

  **cynthiacy** 3 years, 5 months ago

ACF.

F refers to <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/encryption.html>

upvoted 2 times

A company is running an Amazon RDS for PostgreSQL DB instance and wants to migrate it to an Amazon Aurora PostgreSQL DB cluster. The current database is 1 TB in size. The migration needs to have minimal downtime. What is the FASTEST way to accomplish this?

- A. Create an Aurora PostgreSQL DB cluster. Set up replication from the source RDS for PostgreSQL DB instance using AWS DMS to the target DB cluster.
- B. Use the pg_dump and pg_restore utilities to extract and restore the RDS for PostgreSQL DB instance to the Aurora PostgreSQL DB cluster.
- C. Create a database snapshot of the RDS for PostgreSQL DB instance and use this snapshot to create the Aurora PostgreSQL DB cluster.
- D. Migrate data from the RDS for PostgreSQL DB instance to an Aurora PostgreSQL DB cluster using an Aurora Replica. Promote the replica during the cutover.

Suggested Answer: C

Community vote distribution

D (100%)

🗳️ 👤 **[Removed]** Highly Voted 👍 3 years, 8 months ago

D. While C would work, the requirement is minimal downtime.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Migrating.html>

upvoted 14 times

🗳️ 👤 **Pranava_GCP** Most Recent 🕒 1 year, 8 months ago

Selected Answer: D

D. Migrate data from the RDS for PostgreSQL DB instance to an Aurora PostgreSQL DB cluster using an Aurora Replica. Promote the replica during the cutover.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Migrating.html#AuroraPostgreSQL.Migrating.RDSPostgreSQL.Replic>

upvoted 1 times

🗳️ 👤 **IhorK** 1 year, 10 months ago

D

Migrating from RDS PostgreSQL to Aurora PostgreSQL.

You have two options when migrating data from RDS PostgreSQL to Aurora PostgreSQL:

- Use an RDS PostgreSQL snapshot
- Use an Aurora Read Replica for RDS PostgreSQL

<https://aws.amazon.com/blogs/database/migrate-to-an-amazon-aurora-postgresql-instance-from-another-postgresql-source/>

upvoted 1 times

🗳️ 👤 **mraronsimon** 1 year, 10 months ago

Selected Answer: D

The correct answer is D. The keyword is "FASTEST" ~ minimal downtime / minimal effort

"You can migrate your existing Amazon RDS MySQL databases to Amazon Aurora using Aurora Read Replica. This solution is beneficial since it's completely managed and does not involve manually configuring replication functionality to reduce downtime during migration."

"PostgreSQL follows the same process as the one described previously for MySQL for migration."

<https://docs.aws.amazon.com/whitepapers/latest/migrating-databases-to-amazon-aurora/migration-using-aurora-read-replica.html>

upvoted 1 times

🗳️ 👤 **SteveMartin9** 2 years, 4 months ago

Selected Answer: D

Author from the Udemy.com practice test says D is the correct answer.

upvoted 1 times

🗳️ 👤 **sirfans** 2 years, 7 months ago

Selected Answer: D

D is the right option

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: D

C - snapshot will take time (and new transactions lost)

D- Replica is fast

upvoted 2 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: D

D is absolutely correct

upvoted 1 times

🗳️ 👤 **pcpcpc888** 3 years, 3 months ago

We need the SPEEDIEST solution, however, we also need a reasonable solution; how would we deal with database changes when restoring the database when we pick B? AWS perform that job for us when choosing D.

upvoted 2 times

🗳️ 👤 **jove** 3 years, 5 months ago

Keywords : seamless and speediest..

Answer : Option D

upvoted 2 times

🗳️ 👤 **GMartinelli** 3 years, 5 months ago

Selected Answer: D

Option D

upvoted 1 times

🗳️ 👤 **Anuragdba** 3 years, 6 months ago

D: because in question talking about minimum downtime .

upvoted 1 times

🗳️ 👤 **aws4myself** 3 years, 6 months ago

D is the correct answer, Read Replica is faster than snapshots

upvoted 2 times

🗳️ 👤 **guru_ji** 3 years, 6 months ago

D ==>> Correct Answer.

upvoted 1 times

🗳️ 👤 **jayshah7** 3 years, 6 months ago

if we say C is fastest then how to handle changes to source server while snapshot is being transferred and getting applied to target ?

upvoted 2 times

🗳️ 👤 **LMax** 3 years, 7 months ago

Answer D

upvoted 2 times

🗳️ 👤 **jyrajan** 3 years, 7 months ago

Answer D. When you take a snapshot there will be a temporary suspension of I/O Services, now this can be seconds or longer, but our size is 1TB which means that there will be downtime, so based on that the answer is D

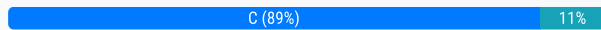
upvoted 1 times

A Database Specialist is migrating a 2 TB Amazon RDS for Oracle DB instance to an RDS for PostgreSQL DB instance using AWS DMS. The source RDS Oracle DB instance is in a VPC in the us-east-1 Region. The target RDS for PostgreSQL DB instance is in a VPC in the use-west-2 Region. Where should the AWS DMS replication instance be placed for the MOST optimal performance?

- A. In the same Region and VPC of the source DB instance
- B. In the same Region and VPC as the target DB instance
- C. In the same VPC and Availability Zone as the target DB instance
- D. In the same VPC and Availability Zone as the source DB instance

Suggested Answer: D

Community vote distribution



zanhsieh 3 years, 8 months ago

C. See the diagram here:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_ReplicationInstance.VPC.html#CHAP_ReplicationInstance.VPC.Configurations.ScenarioVPCPeer

In fact, all the configurations list on above url prefer the replication instance putting into target vpc region / subnet / az.

upvoted 17 times

scottkerker 3 years, 5 months ago

There is also another question related to migration between a RDS resource and a Redshift cluster where they are in different VPCs and the Redshift cluster is the target. It asks readers where should the DMS replication instance be placed and the answer is the replication instance should be placed with the same VPC as the target.

upvoted 2 times

halol 3 years, 8 months ago

D is correct based on :

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_ReplicationInstance.html

it shows Choose the Availability Zone where your source database is located.

upvoted 6 times

BillyMadison 3 years, 8 months ago

Going with D as well since it mentions in your link "reducing load on your source database"

upvoted 1 times

BillyMadison 3 years, 7 months ago

Changing my mind and going with C now. Talked with an AWS solutions architect. He said "target DB instance" has worked better for him in the past.

upvoted 5 times

guru_ji 3 years, 7 months ago

C is correct.

glat mtt kar dena exam mein. All the Best !!

upvoted 1 times

michalf84 7 months, 1 week ago

Selected Answer: B

B there is similar question for redshift migration. You want to reduce latency

upvoted 1 times

MultiAZ 1 year, 4 months ago

Selected Answer: C

It is C - same VPC and AZ as the target.

upvoted 1 times

blah88 1 year, 8 months ago

Selected Answer: B

Answer is B. In the network configuration for database migrations section on this page

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_ReplicationInstance.VPC.html States --> When practical, we recommend that you create a DMS replication instance in the same Region as your target endpoint, and in the same VPC or subnet as your target endpoint.

upvoted 1 times

🗳️ 👤 **Pranava_GCP** 1 year, 8 months ago

Selected Answer: C

C. In the same VPC and Availability Zone as the target DB instance

upvoted 1 times

🗳️ 👤 **aws2023a** 1 year, 9 months ago

https://repost.aws/questions/QUUUL_7NsZS8OBI_GkCRSRRg/why-place-dms-replication-instance-in-target-vpc-instead-of-source-vpc

upvoted 2 times

🗳️ 👤 **leotoras** 1 year, 10 months ago

D is the right answer. It is a heterogeneous migration, some data transformations need to be performed before loading in the target. That's why it should be in the source database VPC.

upvoted 1 times

🗳️ 👤 **f__16** 2 years, 2 months ago

Selected Answer: C

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_ReplicationInstance.VPC.html#CHAP_ReplicationInstance.VPC.Configurations.ScenarioVPCPec

upvoted 1 times

🗳️ 👤 **Vgarg2091** 2 years, 3 months ago

Ans C : Generally, you'd get better performance by placing primary DMS replication instance is in the same AZ as the target DB

upvoted 1 times

🗳️ 👤 **lollyj** 2 years, 5 months ago

Selected Answer: C

Answer: C. I'm new to this site. Is the suggested answer typically the correct answer. I'm really confused. Who provides the suggested answers on this site?

upvoted 2 times

🗳️ 👤 **Zdujgfr567783ff** 2 years, 1 month ago

just a dummy random answer

upvoted 2 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: C

C. In the same VPC and Availability Zone as the target DB instance

(BTW, target region is implicit)

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_ReplicationInstance.VPC.html#CHAP_ReplicationInstance.VPC.Configurations.ScenarioVPCPec

upvoted 3 times

🗳️ 👤 **jnassp1** 3 years, 6 months ago

C - "For best results, we recommend that you locate your replication instance in the same VPC and Availability Zone as your target database, in this case Aurora MySQL." -

upvoted 2 times

🗳️ 👤 **Anuragdba** 3 years, 6 months ago

C: we will get performance improvement if place dms replication instance in target (VPC & availability zone same)

upvoted 1 times

🗳️ 👤 **LB** 3 years, 7 months ago

C is the right answer here ie, the DMS replication instance should be placed in the target AZ.

D would have been correct if the source data needed a lot of filtering and manipulation before transferring data to the target.

upvoted 2 times

🗳️ 👤 **Paulv82003** 1 year, 11 months ago

But data will needed a lot of filtering and manipulation before transferring data to the target going from Oracle to PostgreSQL, D is the correct answer

upvoted 1 times

🗨️ 👤 **aws4myself** 3 years, 7 months ago

C => AZ of target instance

upvoted 1 times

🗨️ 👤 **guru_ji** 3 years, 7 months ago

C ==>> Correct Answer

upvoted 1 times

The Development team recently executed a database script containing several data definition language (DDL) and data manipulation language (DML) statements on an Amazon Aurora MySQL DB cluster. The release accidentally deleted thousands of rows from an important table and broke some application functionality.

This was discovered 4 hours after the release. Upon investigation, a Database Specialist tracked the issue to a DELETE command in the script with an incorrect

WHERE clause filtering the wrong set of rows.

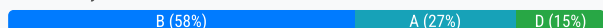
The Aurora DB cluster has Backtrack enabled with an 8-hour backtrack window. The Database Administrator also took a manual snapshot of the DB cluster before the release started. The database needs to be returned to the correct state as quickly as possible to resume full application functionality. Data loss must be minimal.

How can the Database Specialist accomplish this?

- A. Quickly rewind the DB cluster to a point in time before the release using Backtrack.
- B. Perform a point-in-time recovery (PITR) of the DB cluster to a time before the release and copy the deleted rows from the restored database to the original database.
- C. Restore the DB cluster using the manual backup snapshot created before the release and change the application configuration settings to point to the new DB cluster.
- D. Create a clone of the DB cluster with Backtrack enabled. Rewind the cloned cluster to a point in time before the release. Copy deleted rows from the clone to the original database.

Suggested Answer: D

Community vote distribution



edmondme Highly Voted 3 years, 7 months ago

D is right, you can create a clone with backtrack if the database was created with backtrack which it was in this case. It's either B (pitr) or D (backtrack). Backtrack is faster. A is wrong, because if you backtrack, you lose the data that users entered for the past 4 hours. you want to clone to another area and copy the data that was lost.

upvoted 25 times

Jiang_aws1 2 years, 7 months ago

- X A. May lose 4 hrs users entered
- X B. Take too long (restored -> copy delete rows to Org-DB)
- X C. Take too long & May lose 4 hrs users entered
- D. Better a. clone is faster than restore b. Copy delete rows to Org-DB

upvoted 2 times

Jiang_aws1 2 years, 7 months ago

Only A, D using "Rewind" which is good
D is correct if we can rewind "clone with Backtrack enabled" Anyone know ?

upvoted 3 times

aqiao 1 year, 9 months ago

Agree with you, D is the best
option:<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Backtrack.html>
upvoted 1 times

aqiao 1 year, 8 months ago

Correct my answer, it should be B. D is wrong because "you can't backtrack a database clone to a time before that database clone was created" from the backtrack limitation here
:<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Backtrack.html>
upvoted 2 times

lelesp6 2 years ago

D its wrong, da backtrack only right after you create clone.

upvoted 3 times

🗨️ 👤 **halol** Highly Voted 3 years, 8 months ago

A&B correct. however, A cause loss data for last 4 hours but it's quickly will take just few minutes to rewind the database to 4 hours before, however B will keep database and fix it by reinsert deleted records from database which created from restore point in time this will take more time but without loss data as the question refer to
"Data loss must be minimal"
So I guess B the correct answer
upvoted 24 times

🗨️ 👤 **ChauPhan** 3 years, 6 months ago

B is wrong because "Perform a point-in-time recovery (PITR) of the DB cluster to a time before the release"
I example: The release time was 5AM, and you found that the records are deleted at 9AM. So why we recover the database "before 5AM" (release time). ==> we still lose 4 hours.
The correct way to minimize data loss is: recover as nearest as possible time such as 8:55AM) then copy the deletion data. So only 5 mins data loss. Only this minimizes data loss.
upvoted 4 times

🗨️ 👤 **khchan123** 3 years ago

Agree with ChauPhan. B is wrong. A is correct.
upvoted 1 times

🗨️ 👤 **swakan** 2 years, 8 months ago

Option B suggest to restore the db to a point before the records were deleted, then copy the required records from this newly restored db (using PITR) to the original DB. So this way, all the other changes made on the DB are still there, while we copied the deleted rows from the PITR db. I think option B is correct. Please correct if wrong.
upvoted 3 times

🗨️ 👤 **Maze** 2 years, 7 months ago

PITR is restored another cluster with Product cluster. Executing PITR doesn't mean back-forward. we dont lose 4 hours. we can export and import from PITR restored cluster.
upvoted 1 times

🗨️ 👤 **missipssamarsh** Most Recent 1 year, 2 months ago

Selected Answer: B

X D : You can't backtrack a database clone to a time before that database clone was created.
<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Backtrack.html>
upvoted 1 times

🗨️ 👤 **Bezi** 1 year, 2 months ago

I will go for "A". Here's an excerpt from the below link:
<https://aws.amazon.com/blogs/aws/amazon-aurora-backtrack-turn-back-time/>
... After that regrettable moment when all seems lost, you simply pause your application, open up the Aurora Console, select the cluster, and click Backtrack DB cluster. Then you select Backtrack and choose the point in time just before your epic fail, and click Backtrack DB cluster.
upvoted 2 times

🗨️ 👤 **MultiAZ** 1 year, 4 months ago

Selected Answer: B

The answer is B.
With A you would lose all the data gathered in the last 4 hours
upvoted 1 times

🗨️ 👤 **Hisayuki** 1 year, 4 months ago

Selected Answer: A

BackTrack - If you accidentally issue a Delete statement without a Where clause (that is, delete all items) or drop a table, you can use Backtrack to quickly revert to the previous state. You can do the same thing with PITR, but it takes time to restore from an existing backup to a different cluster/instance.
upvoted 1 times

🗨️ 👤 **Hisayuki** 1 year, 4 months ago

Sorry, I noticed the window is 8 hours. We have to use PITR not to lose the data of the 4 hours. The answer is B.
upvoted 1 times

🗨️ 👤 **Jrhp** 1 year, 5 months ago

Selected Answer: A

Backtracking "rewinds" the DB cluster to the time you specify. Backtracking is not a replacement for backing up your DB cluster so that you can restore it to a point in time. However, backtracking provides the following advantages over traditional backup and restore:

You can easily undo mistakes. If you mistakenly perform a destructive action, such as a DELETE without a WHERE clause, you can backtrack the DB cluster to a time before the destructive action with minimal interruption of service

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Backtrack.html#AuroraMySQL.Managing.Backtrack.Overview>
upvoted 1 times

🗳️ 👤 **AmbrishK** 1 year, 6 months ago

Selected Answer: D

This option combines the benefits of both Backtrack and data recovery by cloning the cluster and using Backtrack to recover lost data. After that, you can copy the missing rows back into the original database. It allows for minimal data loss.

Based on the goal of minimizing data loss and ensuring a quick recovery, option D seems like the best choice. It leverages Backtrack, provides a way to recover the deleted data, and then copy it back into the original database. However, option B is also a valid choice if Backtrack is not an option or if the deleted data is outside the backtrack window.

upvoted 1 times

🗳️ 👤 **Germaneli** 1 year, 8 months ago

Selected Answer: A

A is correct, because backtrack is the fastest way to "[return] the database to the correct state as quickly as possible", as the questions asks.

There is no mention or requirement of retaining potentially lost user input.

Additionally, restoring the deleted rows while keeping actual user input from the last 4 hours after the loss might infringe foreign key constraints within the database, making the data inconsistent.

upvoted 2 times

🗳️ 👤 **aqiao** 1 year, 8 months ago

Selected Answer: B

A x: It will cause 4 hours data loss after release

B yes

C x: It will cause 4 hours data loss after release and involve additional change

D x: You can't backtrack a database clone to a time before that database clone was created from the backtrack limitation here

:<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Backtrack.html>

upvoted 2 times

🗳️ 👤 **SamDDD** 1 year, 10 months ago

Selected Answer: A

You can't backtrack a database clone to a time before that database clone was created. However, you can use the original database to backtrack to a time before the clone was created

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Backtrack.html>

upvoted 1 times

🗳️ 👤 **mraronsimon** 1 year, 10 months ago

Selected Answer: B

A & C (incorrect) - dataloss!

D (incorrect) - "You can't backtrack a database clone to a time before that database clone was created. However, you can use the original database to backtrack to a time before the clone was created."

The correct answer is B.

However, D could be better with a little change:

You should use the clone db as new production and you should rewind on the original. Copy deleted rows from original to the clone.

upvoted 2 times

🗳️ 👤 **mraronsimon** 1 year, 10 months ago

Reference: <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Backtrack.html>

upvoted 1 times

🗳️ 👤 **MrAliMohsan** 1 year, 11 months ago

Selected Answer: B

Option D would have been correct if it has suggested to rewind the DB cluster not the cloned cluster. Since you cannot rewind the cloned cluster before the ti

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Backtrack.html#:~:text=You%20can%27t%20backtrack%20a%2>

upvoted 1 times

🗨️ 👤 **aviathor** 1 year, 12 months ago

Selected Answer: B

xA. Rewinding the database means restoring the deleted data, but also undoing all other changes to the database => data loss
B. Perform a point-in-time recovery (PITR) of the DB cluster to a time before the release and copy the deleted rows from the restored database to the original database. => This is the best course of action.
xC. This also entails losing all of the changes made after the deletion event.
xD. This does not work because you cannot rewind a clone to a time before the clone was created.
upvoted 3 times

🗨️ 👤 **aviathor** 1 year, 12 months ago

xA. Rewinding the database means restoring the deleted data, but also undoing all other changes to the database => data loss
B. Perform a point-in-time recovery (PITR) of the DB cluster to a time before the release and copy the deleted rows from the restored database to the original database. => This is the best course of action.
xC. This also entails losing all of the changes made after the deletion event.
xD. This does not work because you cannot rewind a clone to a time before the clone was created.
upvoted 1 times

🗨️ 👤 **backbencher2022** 2 years, 2 months ago

D is right - <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Backtrack.html>
upvoted 1 times

🗨️ 👤 **sk1974** 2 years, 2 months ago

<https://aws.amazon.com/getting-started/hands-on/aurora-cloning-backtracking/> <--- D is the answer
upvoted 1 times

A company is load testing its three-tier production web application deployed with an AWS CloudFormation template on AWS. The Application team is making changes to deploy additional Amazon EC2 and AWS Lambda resources to expand the load testing capacity. A Database Specialist wants to ensure that the changes made by the Application team will not change the Amazon RDS database resources already deployed. Which combination of steps would allow the Database Specialist to accomplish this? (Choose two.)

- A. Review the stack drift before modifying the template
- B. Create and review a change set before applying it
- C. Export the database resources as stack outputs
- D. Define the database resources in a nested stack
- E. Set a stack policy for the database resources

Suggested Answer: AD

Community vote distribution

BE (100%)

🗳️ **chicagomassageseeker** Highly Voted 3 years, 8 months ago

Correct Answer B and E.

B. You need to review the change set before accepting changes.

E. use Stack set policy to prevent updates to resources.

upvoted 22 times

🗳️ **RDSBot** Highly Voted 3 years, 7 months ago

B&E : https://docs.amazonaws.cn/en_us/AWSCloudFormation/latest/UserGuide/best-practices.html#cfn-best-practices-changesets

upvoted 5 times

🗳️ **roymunson** Most Recent 1 year, 8 months ago

Selected Answer: BE

Correct Answer B and E.

upvoted 1 times

🗳️ **cat_of_meerkat** 1 year, 8 months ago

Selected Answer: BE

B and E

upvoted 1 times

🗳️ **lollyj** 2 years, 5 months ago

Selected Answer: BE

Answer BE however I don't answer the A. Can someone explain?

upvoted 1 times

🗳️ **chikorita** 1 year, 8 months ago

CFT stacks drift after update but option A says to check the drift status before even making the update (need to use ChangeSet) which is impossible

upvoted 1 times

🗳️ **sayed** 2 years, 7 months ago

Selected Answer: BE

B and E

upvoted 1 times

🗳️ **KaranGandhi30** 3 years ago

Selected Answer: BE

B. change set will assure that App Dev modification are not impacting resources.

E. Policy rules will be additional to fail deployment if the resources are changed.

upvoted 1 times

🗳️ **novice_expert** 3 years ago

Selected Answer: BE

https://docs.amazonaws.cn/en_us/AWSCloudFormation/latest/UserGuide/best-practices.html#cfn-best-practices-changesets
upvoted 2 times

🗉 👤 **RotterDam** 3 years, 2 months ago
Got this question in my exam. (i cleared it). BE is correct
upvoted 2 times

🗉 👤 **RotterDam** 3 years, 2 months ago
Selected Answer: BE
B & E are the right choices
upvoted 1 times

🗉 👤 **guru_ji** 3 years, 6 months ago
Correct Answer: B and E
upvoted 1 times

🗉 👤 **aws4myself** 3 years, 6 months ago
its BE, stack policy and changesets
upvoted 1 times

🗉 👤 **ChauPhan** 3 years, 7 months ago
B & E are the answers.
A is not clear and headache as we need to review the CF manually, not recommended.
C D won't help
upvoted 3 times

🗉 👤 **LMax** 3 years, 7 months ago
I would go with B & E
upvoted 2 times

🗉 👤 **myutran** 3 years, 7 months ago
Ans: BE
upvoted 1 times

🗉 👤 **Faz** 3 years, 7 months ago
Answer is BE
upvoted 1 times

🗉 👤 **Ashoks** 3 years, 7 months ago
B and E
upvoted 1 times

A manufacturing company's website uses an Amazon Aurora PostgreSQL DB cluster.

Which configurations will result in the LEAST application downtime during a failover? (Choose three.)

- A. Use the provided read and write Aurora endpoints to establish a connection to the Aurora DB cluster.
- B. Create an Amazon CloudWatch alert triggering a restore in another Availability Zone when the primary Aurora DB cluster is unreachable.
- C. Edit and enable Aurora DB cluster cache management in parameter groups.
- D. Set TCP keepalive parameters to a high value.
- E. Set JDBC connection string timeout variables to a low value.
- F. Set Java DNS caching timeouts to a high value.

Suggested Answer: ABC

Community vote distribution

ACE (100%)

 **Mickysingh** Highly Voted 3 years, 8 months ago


Ans is ACE

upvoted 10 times

 **LMax** 3 years, 7 months ago

Agree with ACE

upvoted 3 times

 **im_not_robot** Highly Voted 2 years, 3 months ago

ACE

A: We should use read/cluster endpoints of aurora. If we use other endpoints like instance endpoint, custom endpoint, it might not work correctly after failover.


B: Aurora failover within region is automatic so we don't need cloudwatch alert to do anything

C. Cluster cache management will help to warm up the cache on the new primary instance -> no high load on the new primary instance

D. F. Long dns cache or connection are enemies of failover so we should avoid that -> E is correct

Sorry for my english


upvoted 8 times

 **roymunson** Most Recent 1 year, 8 months ago

Selected Answer: ACE

Correct Answer A,C and E.

upvoted 1 times

 **IhorK** 1 year, 9 months ago

ACE

Regarding C:

Fast recovery after failover with cluster cache management for Aurora PostgreSQL.

With cluster cache management, you set a specific reader DB instance as the failover target. Cluster cache management ensures that the data in the designated reader's cache is kept synchronized with the data in the writer DB instance's cache. The designated reader's cache with prefilled values is known as a warm cache. If a failover occurs, the designated reader uses values in its warm cache immediately when it's promoted to the new writer DB instance. This approach provides your application much better recovery performance.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.cluster-cache-mgmt.html>

upvoted 2 times

 **satishstechie** 2 years, 5 months ago

ACE should be wright

upvoted 1 times

 **rag1482** 2 years, 5 months ago

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.BestPractices.FastFailover.html>

ACE based on above link

upvoted 3 times

🗨️ 👤 **praffuln** 3 years ago

Selected Answer: ACE

Question is discussing about low recovery time at failure...

upvoted 1 times

🗨️ 👤 **novice_expert** 3 years ago

Selected Answer: ACE

A E is correct

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.BestPractices.html>

A. Use the endpoints to cluster.

x B cloudwatch...

x D. Set TCP keepalive parameters to a high value. (keep it low)

E. Set JDBC connection string timeout variables to a low value.

x F. Set Java DNS caching timeouts to a high value. (keep it low)

C is here

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.cluster-cache-mgmt.html>

C. Edit and enable Aurora DB cluster cache management in parameter groups. (keeps cache in sync on replica)

upvoted 3 times

🗨️ 👤 **RotterDam** 3 years, 2 months ago

Got this question in my exam. (i cleared it). ACE is correct

upvoted 5 times

🗨️ 👤 **lokeshM** 3 years, 6 months ago

Ans is ACE

upvoted 1 times

🗨️ 👤 **aws4myself** 3 years, 6 months ago

its ACE

upvoted 2 times

🗨️ 👤 **guru_ji** 3 years, 7 months ago

A,C,E ==>> Correct Answer.

upvoted 2 times

🗨️ 👤 **ChauPhan** 3 years, 7 months ago

A E is correct

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.BestPractices.html>

C is here

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.cluster-cache-mgmt.html>

upvoted 2 times

🗨️ 👤 **Dip11** 3 years, 7 months ago

A and E are correct. Not sure of third one. C is for keeping the cache warm on replica so that after failover performance is not degraded. Purpose of CCM is not to reduce downtime.

upvoted 1 times

🗨️ 👤 **jyrajan** 3 years, 7 months ago

Based on the link below, the answer looks like it should be ADE. Not sure where it says C is an option to ensure a fast failover.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.BestPractices.html#AuroraPostgreSQL.BestPractices.FastFailover.T>

upvoted 1 times

🗨️ 👤 **ChauPhan** 3 years, 7 months ago

D is not correct. You should set TCP keepalive to low value, not high one.



Enabling TCP keepalive parameters and setting them aggressively ensures that if your client is no longer able to connect to the database, then any active are quickly closed. This action allows the application to react appropriately, such as by picking a new host to connect to

upvoted 2 times

🗨️ 👤 **Windy** 3 years, 7 months ago

Answer is ACE

upvoted 1 times

  **myutran** 3 years, 7 months ago

Ans: ACE

upvoted 1 times

A company is hosting critical business data in an Amazon Redshift cluster. Due to the sensitive nature of the data, the cluster is encrypted at rest using AWS

KMS. As a part of disaster recovery requirements, the company needs to copy the Amazon Redshift snapshots to another Region.

Which steps should be taken in the AWS Management Console to meet the disaster recovery requirements?

- A. Create a new KMS customer master key in the source Region. Switch to the destination Region, enable Amazon Redshift cross-Region snapshots, and use the KMS key of the source Region.
- B. Create a new IAM role with access to the KMS key. Enable Amazon Redshift cross-Region replication using the new IAM role, and use the KMS key of the source Region.
- C. Enable Amazon Redshift cross-Region snapshots in the source Region, and create a snapshot copy grant and use a KMS key in the destination Region.
- D. Create a new KMS customer master key in the destination Region and create a new IAM role with access to the new KMS key. Enable Amazon Redshift cross-Region replication in the source Region and use the KMS key of the destination Region.

Suggested Answer: A

Reference:

<https://docs.aws.amazon.com/redshift/latest/mgmt/working-with-snapshots.html>

The screenshot shows the AWS Documentation page for Amazon Redshift Cluster Management Guide. The breadcrumb trail is: AWS > Documentation > Amazon Redshift > Cluster Management Guide. The page title is 'Managing snapshots using the AWS CLI and Amazon Redshift API'. The left sidebar shows a table of contents with 'Snapshots' expanded, listing 'Managing snapshots using the console', 'Managing snapshots using the AWS SDK for Java', and 'Managing snapshots using the AWS CLI and Amazon Redshift API'. The main content area is titled 'Overview' and contains the following text: 'Snapshots are point-in-time backups of a cluster. There are two types of snapshots: *automated* and *manual*. Amazon Redshift stores these snapshots internally in Amazon S3 by using an encrypted Secure Sockets Layer (SSL) connection. Amazon Redshift automatically takes incremental snapshots that track changes to the cluster since the previous automated snapshot. Automated snapshots retain all of the data required to restore a cluster from a snapshot. You can create a snapshot schedule to control when automated snapshots are taken, or you can take a manual snapshot any time. When you restore from a snapshot, Amazon Redshift creates a new...'. At the bottom, a 'Community vote distribution' bar shows three options: C (80%), 13%, and 7%.

PietraOra Highly Voted 3 years, 6 months ago

I think C

If you want to enable cross-Region snapshot copy for an AWS KMS–encrypted cluster, you must configure a snapshot copy grant for a root key in the destination AWS Region

Source-Region : configure a cross-Region snapshot for an AWS KMS–encrypted cluster

In Destination AWS Region : choose the AWS Region to which to copy snapshots.

<https://docs.aws.amazon.com/redshift/latest/mgmt/managing-snapshots-console.html#xregioncopy-kms-encrypted-snapshot>

upvoted 6 times

Pazzooo Most Recent 1 year, 4 months ago

C

See <https://aws.amazon.com/blogs/big-data/migrate-your-amazon-redshift-cluster-to-another-aws-region/>

upvoted 1 times

Hisayuki 1 year, 4 months ago

Selected Answer: C

For encrypted snapshots - configure cross region snapshots and additionally specify a snapshot copy grant which requires a KMS key

upvoted 1 times

AmbrishK 1 year, 6 months ago

Selected Answer: C

C. Enable Amazon Redshift cross-Region snapshots in the source Region, and create a snapshot copy grant and use a KMS key in the destination Region.

Here's why this option is the correct choice:

Enabling cross-Region snapshots in the source Region is necessary to initiate the snapshot copying process.

Creating a snapshot copy grant allows you to define permissions and configurations for copying snapshots to the destination Region. It is an essential step in setting up snapshot replication.

Using a KMS key in the destination Region ensures that the copied snapshots are encrypted with a key specific to that Region. This maintains data security during replication.

upvoted 1 times

  **roymunson** 1 year, 8 months ago

Selected Answer: C


It's C:

<https://docs.aws.amazon.com/redshift/latest/mgmt/working-with-db-encryption.html#configure-snapshot-copy-grant>

It can't be (A) because when you read the docs it is written "[...] 2. In the SOURCE AWS Region, enable copying of snapshots and ...".

B & D is about replication and not copying a snapshot.

upvoted 1 times

  **Germaneli** 1 year, 8 months ago

Selected Answer: A

The question is about cross-region snapshot copy (read carefully), not about cross-region replication, so B and D are out.

From the remaining A and C, I would tend to A because the KMS key is needed from the source, the target only needs a grant on it.

upvoted 1 times

  **aqiao** 1 year, 8 months ago


Selected Answer: C

Seek "Copying AWS KMS-encrypted snapshots to another AWS Region" from

<https://docs.aws.amazon.com/redshift/latest/mgmt/working-with-db-encryption.html>

No need to create IAM role

upvoted 1 times

  **mraronsimon** 1 year, 10 months ago

Selected Answer: C

"To copy snapshots for AWS KMS-encrypted clusters to another AWS Region, create a grant for Amazon Redshift to use a customer managed key in the destination AWS Region. Then choose that grant when you enable copying of snapshots in the source AWS Region."

Reference: <https://docs.aws.amazon.com/redshift/latest/mgmt/working-with-snapshots.html>

upvoted 1 times

  **adelcold** 1 year, 11 months ago

Selected Answer: D

<https://docs.aws.amazon.com/redshift/latest/mgmt/working-with-db-encryption.html#configure-snapshot-copy-grant>

upvoted 1 times

  **adelcold** 1 year, 11 months ago

D is correct

<https://docs.aws.amazon.com/redshift/latest/mgmt/working-with-db-encryption.html#configure-snapshot-copy-grant>

upvoted 1 times

  **ken_test1234** 2 years, 2 months ago

Selected Answer: C

because of this documentation <https://docs.aws.amazon.com/redshift/latest/mgmt/working-with-snapshots.html>



Copying snapshots to another AWS Region article shows c is the answer

upvoted 1 times

  **Mintwater** 2 years, 1 month ago

Why is D not correct?

upvoted 1 times

  **Sathish_dbs** 1 year, 7 months ago

because there is no cross region replication with redshift

upvoted 1 times

  **teo2157** 2 years, 4 months ago

Selected Answer: C

The answer is explained here: <https://docs.aws.amazon.com/redshift/latest/mgmt/working-with-db-encryption.html#working-with-aws-kms>, look for "copying AWS KMS–encrypted snapshots to another AWS Region"

upvoted 1 times

  **lollyj** 2 years, 5 months ago

Selected Answer: D

I thought keys are region specific and one will need to be created in the destination region

upvoted 1 times

  **Dantas** 2 years, 12 months ago

Selected Answer: C

<https://docs.aws.amazon.com/redshift/latest/mgmt/working-with-db-encryption.html>

upvoted 2 times

  **novice_expert** 3 years ago

Selected Answer: C

-A,B,D are incorrect

-C

new KMS in the destination Region

-> snapshot copy grant in the destination Region specifying the new key



->In the source Region, configure cross-Region snapshots for the Amazon Redshift cluster specifying

- the destination Region,

- the snapshot copy grant,

- and retention periods for the snapshot.

upvoted 1 times

  **victornj** 3 years, 4 months ago

Question is poorly written . In absence of true answer C is right. It is right because A,B,D are not correct. You definitely need snapshot copy grant in destination region but based on that region key. Answer C does not say which key.

upvoted 3 times

  **user0001** 3 years, 2 months ago

agree with you

upvoted 1 times

  **MahiShai** 3 years, 4 months ago

C is correct ans

upvoted 2 times

A company has a production Amazon Aurora Db cluster that serves both online transaction processing (OLTP) transactions and compute-intensive reports. The reports run for 10% of the total cluster uptime while the OLTP transactions run all the time. The company has benchmarked its workload and determined that a six- node Aurora DB cluster is appropriate for the peak workload.

The company is now looking at cutting costs for this DB cluster, but needs to have a sufficient number of nodes in the cluster to support the workload at different times. The workload has not changed since the previous benchmarking exercise.

How can a Database Specialist address these requirements with minimal user involvement?

- A. Split up the DB cluster into two different clusters: one for OLTP and the other for reporting. Monitor and set up replication between the two clusters to keep data consistent.
- B. Review all evaluate the peak combined workload. Ensure that utilization of the DB cluster node is at an acceptable level. Adjust the number of instances, if necessary.
- C. Use the stop cluster functionality to stop all the nodes of the DB cluster during times of minimal workload. The cluster can be restarted again depending on the workload at the time.
- D. Set up automatic scaling on the DB cluster. This will allow the number of reader nodes to adjust automatically to the reporting workload, when needed.


Suggested Answer: D

Community vote distribution

D (100%)

  **chicagomassageseeker** Highly Voted 3 years, 8 months ago

Answer D. You dont setup a separete cluster for reporting. Aurora does OLTP and reporting from same cluster. just use autoscaling.
upvoted 15 times

  **MultiAZ** Most Recent 1 year, 4 months ago

Selected Answer: D

Out of the provided options, D seems to be the best.
upvoted 1 times

  **Pranava_GCP** 1 year, 8 months ago

Selected Answer: D

D. Set up Auto Scaling --> minimal user involvement

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Integrating.AutoScaling.html#Aurora.Integrating.AutoScaling.Concepts>

"Aurora Auto Scaling enables your Aurora DB cluster to handle sudden increases in connectivity or workload. When the connectivity or workload decreases, Aurora Auto Scaling removes unnecessary Aurora Replicas so that you don't pay for unused provisioned DB instances."

upvoted 2 times

  **Zimboguru** 2 years, 8 months ago



Selected Answer: D

D is correct
upvoted 1 times

  **novice_expert** 3 years ago

Selected Answer: D

Auto-scaling will scale-up/down number of read-replicas
upvoted 1 times

  **guru_ji** 3 years, 6 months ago

D ==>> Correct Answer.
upvoted 1 times

  **aws4myself** 3 years, 6 months ago

D => autoscaling on Aurora cluster
upvoted 1 times

🗨️ 👤 **guru_ji** 3 years, 6 months ago

D ==>> Correct Answer.

upvoted 1 times

🗨️ 👤 **LMax** 3 years, 7 months ago

D for sure.

upvoted 2 times

🗨️ 👤 **myutran** 3 years, 7 months ago

Ans: D

upvoted 1 times

🗨️ 👤 **JobinAkaJoe** 3 years, 7 months ago

D is the correct answer. Auto-scaling will scale-up/down number of read-replicas automatically based on the workload.

upvoted 1 times

🗨️ 👤 **Ebi** 3 years, 7 months ago

Answer is D

upvoted 1 times

🗨️ 👤 **BillyC** 3 years, 7 months ago

D is Correct

upvoted 2 times

A company is running a finance application on an Amazon RDS for MySQL DB instance. The application is governed by multiple financial regulatory agencies.

The RDS DB instance is set up with security groups to allow access to certain Amazon EC2 servers only. AWS KMS is used for encryption at rest. Which step will provide additional security?

- A. Set up NACLs that allow the entire EC2 subnet to access the DB instance
- B. Disable the master user account
- C. Set up a security group that blocks SSH to the DB instance
- D. Set up RDS to use SSL for data in transit

Suggested Answer: D

Reference:

<https://aws.amazon.com/blogs/database/applying-best-practices-for-securing-sensitive-data-in-amazon-rds/>

AWS Database Blog

Applying best practices for securing sensitive data in Amazon RDS

by Syed Jaffry | on 03 APR 2019 | in [Amazon RDS](#), [Database](#) | [Permalink](#) | [Comments](#) | [Share](#)

In the [first post](#) of the series, I described some generic security concepts and corresponding AWS security controls that can be applied to data stores on AWS. Using these, you can create a stronger security posture around your data. In this second post, I demonstrate how these concepts can be implemented to [Amazon RDS](#) databases.

Although many of the implementation examples are common to all RDS database engines, a few might differ based on the individual engine type. In these cases, I include implementation examples from [Amazon Aurora](#) with MySQL compatibility but also point you to where to get the information for other database engines.

Let's walk through the implementation of the security concepts in the order in which they were described in the first post.

Community vote distribution

D (100%)

jove Highly Voted 3 years, 6 months ago

Selected Answer: D

This needs to be corrected in the question: Instead of "AWS KMS is used to encrypt data in transit" it should be "AWS KMS is used to encrypt data at rest".

upvoted 18 times

Pranava_GCP Most Recent 1 year, 8 months ago

Selected Answer: D

D. Set up RDS to use SSL for data in transit

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.SSL.html>

"SSL/TLS connections provide a layer of security by encrypting data that moves between your client and DB instance. Optionally, your SSL/TLS connection can perform server identity verification by validating the server certificate installed on your DB instance."

upvoted 2 times

megramlak 2 years ago

D is the correct answer, in AWS RDS we cannot SSH into underlying instances

upvoted 2 times

🗨️ 👤 **im_not_robot** 2 years, 3 months ago

Selected Answer: D

A is incorrect since it allows all ec2 instances can connect to db

B is incorrect because it doesn't help security

C is incorrect because security group doesn't have 'Deny' rule

D is correct because the data is encrypted in transit

upvoted 1 times

🗨️ 👤 **vkruger** 2 years, 6 months ago

Selected Answer: D

Correct answer is D

upvoted 1 times

🗨️ 👤 **novice_expert** 3 years ago

Selected Answer: D

SSL for data in transit

upvoted 1 times

🗨️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: D

After text correct in Q, D is good for security in transit

upvoted 3 times

🗨️ 👤 **Sp230** 3 years, 6 months ago

D set up ssl

upvoted 3 times

A company needs a data warehouse solution that keeps data in a consistent, highly structured format. The company requires fast responses for end-user queries when looking at data from the current year, and users must have access to the full 15-year dataset, when needed. This solution also needs to handle a fluctuating number incoming queries. Storage costs for the 100 TB of data must be kept low.



Which solution meets these requirements?

- A. Leverage an Amazon Redshift data warehouse solution using a dense storage instance type while keeping all the data on local Amazon Redshift storage. Provision enough instances to support high demand.
- B. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. Provision enough instances to support high demand.
- C. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. Enable Amazon Redshift Concurrency Scaling.
- D. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. Leverage Amazon Redshift elastic resize.

Suggested Answer: C

Community vote distribution

C (100%)

  **BillyMadison** Highly Voted 3 years, 8 months ago

C.

<https://docs.aws.amazon.com/redshift/latest/dg/concurrency-scaling.html>

"ith the Concurrency Scaling feature, you can support virtually unlimited concurrent users and concurrent queries, with consistently fast query performance. When concurrency scaling is enabled, Amazon Redshift automatically adds additional cluster capacity when you need it to process an increase in concurrent read queries. Write operations continue as normal on your main cluster. Users always see the most current data, whether the queries run on the main cluster or on a concurrency scaling cluster. You're charged for concurrency scaling clusters only for the time they're in use. For more information about pricing, see Amazon Redshift pricing. You manage which queries are sent to the concurrency scaling cluster by configuring WLM queues. When you enable concurrency scaling for a queue, eligible queries are sent to the concurrency scaling cluster instead of waiting in line."

upvoted 14 times

  **aviathor** Most Recent 1 year, 12 months ago

These days, wouldn't you rather use RA3 nodes that automatically unload unused data to 23

upvoted 2 times

  **novice_expert** 3 years ago

Selected Answer: C

<https://docs.aws.amazon.com/redshift/latest/dg/concurrency-scaling.html>


upvoted 1 times

  **JustPassIt** 3 years, 2 months ago

D.

Main cluster can't be a single-node cluster.



upvoted 1 times

  **GMartinelli** 3 years, 5 months ago

Selected Answer: C



Option C

upvoted 2 times

  **GaryY** 3 years, 7 months ago

Why not D? I think D (Redshift elastic resize) is also correct.

upvoted 1 times

  **surisits** 3 years, 7 months ago

Elastic resize is only available for Amazon Redshift clusters that use the EC2-VPC platform. Option of "Dense Storage" and "Elastic resize" is mutually exclusive I think

upvoted 1 times

🗨️ 👤 **Jiang_aws1** 2 years, 7 months ago

<https://aws.amazon.com/redshift/faqs/>

Q: What is Elastic Resize and how is it different from Concurrency Scaling?

Elastic Resize adds or removes nodes from a single Redshift cluster within minutes to manage its query throughput. For example, an ETL workload for certain hours in a day or month-end reporting may need additional Amazon Redshift resources to complete on time. Concurrency Scaling adds additional cluster resources to increase the overall query concurrency.

upvoted 3 times

🗨️ 👤 **GMartinelli** 3 years, 7 months ago

C. But I would argue that you could use redshift spectrum for everything, I just don't know if you can use it on S3 IA for the historical data, but it might work.

upvoted 1 times

🗨️ 👤 **im_not_robot** 2 years, 3 months ago

redshift spectrum doesn't have fast response so C is incorrect.

upvoted 1 times

🗨️ 👤 **guru_ji** 3 years, 7 months ago

C ==>> Correct

upvoted 2 times

🗨️ 👤 **aws4myself** 3 years, 7 months ago

C => S3 + local storage + concurrency scaling

upvoted 1 times

🗨️ 👤 **LMax** 3 years, 7 months ago

Agree with C

upvoted 1 times

🗨️ 👤 **myutran** 3 years, 7 months ago

Ans: C

upvoted 1 times

🗨️ 👤 **Ebi** 3 years, 7 months ago

C is the answer

upvoted 2 times

🗨️ 👤 **BillyC** 3 years, 7 months ago

C Here

upvoted 2 times

A gaming company wants to deploy a game in multiple Regions. The company plans to save local high scores in Amazon DynamoDB tables in each Region. A

Database Specialist needs to design a solution to automate the deployment of the database with identical configurations in additional Regions, as needed. The solution should also automate configuration changes across all Regions.

Which solution would meet these requirements and deploy the DynamoDB tables?

- A. Create an AWS CLI command to deploy the DynamoDB table to all the Regions and save it for future deployments.
- B. Create an AWS CloudFormation template and deploy the template to all the Regions.
- C. Create an AWS CloudFormation template and use a stack set to deploy the template to all the Regions.
- D. Create DynamoDB tables using the AWS Management Console in all the Regions and create a step-by-step guide for future deployments.

Suggested Answer: B

Community vote distribution

C (100%)

🗲️ 👤 **helpaws** Highly Voted 3 years, 7 months ago

C here

upvoted 11 times

🗲️ 👤 **user0001** 3 years, 2 months ago

true but wording is wrong as well

upvoted 1 times

🗲️ 👤 **Anuragdba** Highly Voted 3 years, 6 months ago

C : Use CloudFormation StackSets to Provision Resources Across Multiple AWS Accounts and Regions

<https://aws.amazon.com/blogs/aws/use-cloudformation-stacksets-to-provision-resources-across-multiple-aws-accounts-and-regions/>

upvoted 5 times

🗲️ 👤 **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: C

C. Create an AWS CloudFormation template and use a stack set to deploy the template to all the Regions.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/what-is-cfnstacksets.html>

"AWS CloudFormation StackSets extends the capability of stacks by enabling you to create, update, or delete stacks across multiple accounts and AWS Regions with a single operation. Using an administrator account, you define and manage an AWS CloudFormation template, and use the template as the basis for provisioning stacks into selected target accounts across specified AWS Regions."

upvoted 2 times

🗲️ 👤 **adelcold** 1 year, 11 months ago

Selected Answer: C

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/what-is-cfnstacksets.html>

upvoted 2 times

🗲️ 👤 **f__16** 2 years, 2 months ago

Selected Answer: C

Stack sets are used to deploy to multiple regions.

upvoted 3 times

🗲️ 👤 **Dean791** 2 years, 2 months ago

Selected Answer: C

C is the best choice

A & D so manual

upvoted 1 times

🗲️ 👤 **lollyj** 2 years, 5 months ago

Selected Answer: C

Stack sets to deploy to regions

upvoted 1 times

🗨️ **sayed** 2 years, 7 months ago

Selected Answer: C

C

automate the deployment of the database with identical configurations in additional Regions (leads us to cloud formation), as needed. The solution should also automate configuration changes across all Regions (leads us to stack set)

upvoted 1 times

🗨️ **novice_expert** 3 years ago

Selected Answer: C

Create an AWS CloudFormation template

-> use a stack set

-> deploy the template to all the Regions.

upvoted 1 times

🗨️ **GMartinelli** 3 years, 5 months ago

Selected Answer: C

Option C

upvoted 2 times

🗨️ **astood** 3 years, 6 months ago

C >>> Absolutely agreed

upvoted 1 times

🗨️ **aws4myself** 3 years, 6 months ago

C => stack set to deploy in multiple regions

upvoted 1 times

🗨️ **guru_ji** 3 years, 7 months ago

Correct Answer ==>> C

upvoted 1 times

🗨️ **ChauPhan** 3 years, 7 months ago

C.

A stack set lets you create stacks in AWS accounts across regions by using a single CloudFormation template

upvoted 2 times

🗨️ **manan728** 3 years, 7 months ago

Yup C it is.

upvoted 1 times

🗨️ **LMax** 3 years, 7 months ago

Answer C

upvoted 2 times

🗨️ **myutran** 3 years, 7 months ago

Ans: C

upvoted 1 times

A team of Database Specialists is currently investigating performance issues on an Amazon RDS for MySQL DB instance and is reviewing related metrics. The team wants to narrow the possibilities down to specific database wait events to better understand the situation. How can the Database Specialists accomplish this?

- A. Enable the option to push all database logs to Amazon CloudWatch for advanced analysis
- B. Create appropriate Amazon CloudWatch dashboards to contain specific periods of time
- C. Enable Amazon RDS Performance Insights and review the appropriate dashboard
- D. Enable Enhanced Monitoring with the appropriate settings

Suggested Answer: C

Community vote distribution

C (85%)

B (15%)

 **BillyC** Highly Voted 3 years, 8 months ago

Yes, C its correct!
upvoted 13 times

 **novice_expert** Highly Voted 3 years ago

Selected Answer: C

C. Enable Amazon RDS Performance Insights will give you metric specifically for database with the dashboard

D. Enhanced monitoring is for OS info
upvoted 5 times

 **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: C

C. Enable Amazon RDS Performance Insights and review the appropriate dashboard

<https://aws.amazon.com/rds/performance-insights/>
upvoted 3 times


 **renfdo** 2 years, 4 months ago

Answer: C
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PerfInsights.UsingDashboard.AnalyzeDBLoad.html
upvoted 2 times


 **SachinGoel** 2 years, 4 months ago

Selected Answer: C

Answer - C
upvoted 1 times

 **sachin** 2 years, 11 months ago

My take is B
upvoted 1 times


 **Balki** 2 years, 11 months ago

Selected Answer: B

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PerfInsights.UsingDashboard.AnalyzeDBLoad.html
upvoted 2 times

 **renfdo** 2 years, 4 months ago

This link, confirm that it's Performance Insights.
Answer: C
upvoted 1 times

 **tugboat** 3 years, 3 months ago

Selected Answer: C

Performance Insights will give you required info

upvoted 2 times

🗉 **astood** 3 years, 6 months ago

C>>> Enable Amazon RDS Performance Insights will give you metric specifically for database with the dashboard to navigate visually around to see database performance activities

upvoted 1 times

🗉 **aws4myself** 3 years, 7 months ago

C is correct

upvoted 2 times

🗉 **guru_ji** 3 years, 7 months ago

Correct Answer ==>> C

upvoted 2 times

🗉 **damaldon** 3 years, 7 months ago

When you use RDS Performance Insights, you can visualize the database load and filter the load by waits, SQL statements, hosts, or users. This way, you can identify which queries are causing issues and view the wait type and wait events associated to that query.

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-mysql-db-performance/>

upvoted 1 times

🗉 **dcabib** 3 years, 7 months ago

Amazon RDS does not gave Fault injection - Just Aurora has this feature ... Answer is B.

upvoted 1 times

🗉 **LMax** 3 years, 7 months ago

Answer C

upvoted 2 times

🗉 **myutran** 3 years, 7 months ago

Ans: C

upvoted 1 times

🗉 **JobinAkaJoe** 3 years, 7 months ago

C, no doubt

upvoted 2 times

🗉 **BillyMadison** 3 years, 8 months ago

C as well.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PerfInsights.Enabling.html

<https://aws.amazon.com/rds/performance-insights/>

<https://aws.amazon.com/blogs/database/tuning-amazon-rds-for-mysql-with-performance-insights/>

upvoted 3 times

A large company is using an Amazon RDS for Oracle Multi-AZ DB instance with a Java application. As a part of its disaster recovery annual testing, the company would like to simulate an Availability Zone failure and record how the application reacts during the DB instance failover activity. The company does not want to make any code changes for this activity.

What should the company do to achieve this in the shortest amount of time?

- A. Use a blue-green deployment with a complete application-level failover test
- B. Use the RDS console to reboot the DB instance by choosing the option to reboot with failover
- C. Use RDS fault injection queries to simulate the primary node failure
- D. Add a rule to the NACL to deny all traffic on the subnets associated with a single Availability Zone

Suggested Answer: C

Reference:

https://wellarchitectedlabs.com/Reliability/300_Testing_for_Resiliency_of_EC2_RDS_and_S3/Lab_Guide.html

The screenshot shows the Well-Architected Labs website. On the left, a sidebar menu lists categories: Operational Excellence, Security, Reliability, and Health Checks & Dependencies. Under Reliability, there are links for 100 Labs, 200 Labs, and 300 Labs. The 'Test Resiliency EC2, RDS, & AZ' link is highlighted. Below this, a '1. Deploy Application' button is visible. On the right, the 'Introduction' section is displayed, explaining the purpose of the lab: to teach the fundamentals of using tests to ensure implementation is resilient to failure by injecting failure modes. It mentions Failure Mode Engineering Analysis (FMEA) and Chaos Engineering. A community vote distribution bar at the bottom shows 'B (93%)' in blue and '7%' in green.

toppic26 Highly Voted 3 years, 6 months ago

B is the correct option. Use the RDS fault injection query to simulate the primary instance failure - This is a trick option. You can test the fault tolerance of your Aurora DB cluster by using fault injection queries. Fault injection queries are issued as SQL commands to an Aurora instance.

Exam Alert:

Fault injection queries can be only used with Aurora DB cluster and NOT with an RDS DB cluster.

upvoted 18 times

MultiAZ Most Recent 1 year, 4 months ago

Selected Answer: B

Answer is B.

Oracle does not support fault injection queries

upvoted 1 times

Pranava_GCP 1 year, 8 months ago

Selected Answer: B

B. Use the RDS console to reboot the DB instance by choosing the option to reboot with failover

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RebootInstance.html

"Rebooting with failover is beneficial when you want to simulate a failure of a DB instance for testing, or restore operations to the original AZ after a failover occurs"

upvoted 3 times

Germaneli 1 year, 8 months ago

Selected Answer: D

A is irrelevant (no software-related change planned).

B is possible, but the requirement mentions a "simulation", not actually doing it.

C - fault injection is feasible only on Aurora, not on RDS for Oracle.

D fulfills the requirement, which is to *simulate* an AZ failure (not only a failover of the database). So I'd vote for that.

upvoted 1 times

🗳️ 👤 **mraronsimon** 1 year, 10 months ago

Selected Answer: B

Exam Alert:

Fault injection queries can be only used with Aurora DB cluster and NOT with an RDS DB cluster.

upvoted 1 times

🗳️ 👤 **Pankaj24hrs** 2 years ago

C - As question as to simulate AZ not doing it manually.

upvoted 1 times

🗳️ 👤 **ken_test1234** 2 years, 2 months ago

Selected Answer: B

Because the fault injection queries will take too much time to fail the engine

upvoted 1 times

🗳️ 👤 **im_not_robot** 2 years, 3 months ago

Selected Answer: B

A is not relevant.

B is correct because it is fast to do

C is incorrect because it is slower than B

D is incorrect because it affects the whole subnets with multiple resources like ec2, alb,...

upvoted 1 times

🗳️ 👤 **lollyj** 2 years, 5 months ago

Selected Answer: B

B is the correct option.

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: B

reboot with failover option.

You can force a crash of an Amazon Aurora instance using the ALTER SYSTEM CRASH fault injection query.

upvoted 3 times

🗳️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: B

Test for failover is done via reboot with failover

upvoted 3 times

🗳️ 👤 **Shunpin** 3 years, 5 months ago

I prefer B.

My point is that you can't simulate real application behaviors. For example in JDBC connection, if the TCP socket disconnect detection is not in the java coder. Your connection is stuck while failover.

upvoted 3 times

🗳️ 👤 **jove** 3 years, 6 months ago

It seems the RDS fault injection queries are supported only on Aurora :

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.FaultInjectionQueries.html>

So, for testing Oracle, B should be the correct option.

upvoted 4 times

🗳️ 👤 **leunamE** 3 years, 6 months ago

Yes, but even if it's Aurora instead of RDS, I think the key is that the question says they want to test how the application behaves during DB instance failover activity. in this case it is always the reboot with failover option.

upvoted 1 times

🗳️ 👤 **2025flakyt** 3 years, 5 months ago

Although B is the correct answer for HA testing but RDS also support fault injection

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/disaster-recovery-resiliency.html>

https://wellarchitectedlabs.com/reliability/300_labs/300_testing_for_resiliency_of_ec2_rds_and_s3/

upvoted 1 times

  **Germaneli** 1 year, 8 months ago

RDS for Oracle doesn't support fault injection, and on the provided links this is also not confirmed.

upvoted 1 times

  **leunamE** 3 years, 6 months ago

Option B.

upvoted 3 times

A company maintains several databases using Amazon RDS for MySQL and PostgreSQL. Each RDS database generates log files with retention periods set to their default values. The company has now mandated that database logs be maintained for up to 90 days in a centralized repository to facilitate real-time and after-the-fact analyses.

What should a Database Specialist do to meet these requirements with minimal effort?

- A. Create an AWS Lambda function to pull logs from the RDS databases and consolidate the log files in an Amazon S3 bucket. Set a lifecycle policy to expire the objects after 90 days.
- B. Modify the RDS databases to publish log to Amazon CloudWatch Logs. Change the log retention policy for each log group to expire the events after 90 days.
- C. Write a stored procedure in each RDS database to download the logs and consolidate the log files in an Amazon S3 bucket. Set a lifecycle policy to expire the objects after 90 days.
- D. Create an AWS Lambda function to download the logs from the RDS databases and publish the logs to Amazon CloudWatch Logs. Change the log retention policy for the log group to expire the events after 90 days.

Suggested Answer: A

Community vote distribution

B (92%)

8%

 **learnaws**  3 years, 8 months ago

I'll go with B because it facilitate real-time
upvoted 12 times

 **szmulder** 3 years, 8 months ago

I think it's due to the question need minimal effort. A need write a lambda is not minimal effort
upvoted 4 times

 **user0001** 3 years, 2 months ago

regardless of the effort, why do you want to write code? (option A).
you should always try to go with built-in fuctionality
upvoted 1 times

 **cloud4gr8** 3 years, 7 months ago

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_LogAccess.html
upvoted 3 times

 **Pranava_GCP**  1 year, 8 months ago

Selected Answer: B

B. Modify the RDS databases to publish log to Amazon CloudWatch Logs. Change the log retention policy for each log group to expire the events after 90 days.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_LogAccess.Procedural.UploadtoCloudWatch.html

"In an on-premises database, the database logs reside on the file system. Amazon RDS doesn't provide host access to the database logs on the file system of your DB instance. For this reason, Amazon RDS lets you export database logs to Amazon CloudWatch Logs. With CloudWatch Logs, you can perform real-time analysis of the log data. You can also store the data in highly durable storage and manage the data with the CloudWatch Logs Agent. "

upvoted 2 times

 **mraronsimon** 1 year, 10 months ago

Selected Answer: B

The answer is B.

"In an on-premises database, the database logs reside on the file system. Amazon RDS doesn't provide host access to the database logs on the file system of your DB instance. For this reason, Amazon RDS lets you export database logs to Amazon CloudWatch Logs. With CloudWatch Logs, you can perform real-time analysis of the log data. You can also store the data in highly durable storage and manage the data with the CloudWatch Logs

Agent."

Reference: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_LogAccess.Procedural.UploadtoCloudWatch.html

upvoted 3 times

🗳️ 👤 **ken_test1234** 2 years, 2 months ago

Selected Answer: B

Because you need a centralized place for analysis and real time data , using s3 will require to search and view the logs 1 by 1 which is not a real time analyses

upvoted 1 times

🗳️ 👤 **Mintwater** 2 years, 2 months ago

B.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_LogAccess.Procedural.UploadtoCloudWatch.html

Publish logs to CloudWatch

upvoted 1 times

🗳️ 👤 **awsjjj** 2 years, 7 months ago

Selected Answer: B

B it is

upvoted 1 times

🗳️ 👤 **sachin** 2 years, 11 months ago

with the least amount of work possible ; _the ans seems to be B but

with have central repository to store logs for post-mortem analysis .. the Ans A seems to be correct..

This is confusing question :)

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: A

no for B because: Logfiles are already created, and we do not want to mix that info with other log entries on CloudWatch

so A:

Lambda function to pull logs from the RDS databases -> consolidate the log files in an Amazon S3 bucket. Set a lifecycle policy to expire the objects after 90 days.

upvoted 1 times

🗳️ 👤 **DevoteamAnalytix** 2 years, 11 months ago

Why not mix logs on CloudWatch? We can use different Log groups and we can filter logs....

upvoted 2 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: B

B is the correction Option

upvoted 1 times

🗳️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: B

B can be scripted too

upvoted 1 times

🗳️ 👤 **soyyodario** 3 years, 3 months ago

Selected Answer: B

B, best option with minimal effort and real time.

upvoted 2 times

🗳️ 👤 **awsmonster** 3 years, 4 months ago

Going with B

upvoted 1 times

🗳️ 👤 **mnzsql365** 3 years, 5 months ago

B is the right ans.


https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_LogAccess.Procedural.UploadtoCloudWatch.html

upvoted 3 times

🗳️ 👤 **andy909** 3 years, 6 months ago



Answer: B

upvoted 1 times

  **GMartinelli** 3 years, 6 months ago

Its a tricky question, the correct option goes down to which of the options uses minimal effort, creating a lambda function or modifying all the databases. The question doesn't tell us how many RDS, so I would choose B

upvoted 1 times

  **guru_ji** 3 years, 6 months ago

B is Correct

upvoted 1 times

  **ChauPhan** 3 years, 6 months ago

Vote B for minimal effort.

upvoted 1 times

A Database Specialist is setting up a new Amazon Aurora DB cluster with one primary instance and three Aurora Replicas for a highly intensive, business-critical application. The Aurora DB cluster has one medium-sized primary instance, one large-sized replica, and two medium sized replicas. The Database Specialist did not assign a promotion tier to the replicas.

In the event of a primary failure, what will occur?

- A. Aurora will promote an Aurora Replica that is of the same size as the primary instance
- B. Aurora will promote an arbitrary Aurora Replica
- C. Aurora will promote the largest-sized Aurora Replica
- D. Aurora will not promote an Aurora Replica

Suggested Answer: A

Reference:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-ug.pdf>

Community vote distribution

C (100%)

🗳️ 👤 **pan24** Highly Voted 3 years, 8 months ago

ANS: C

Priority: If you don't select a value, the default is tier-1. This priority determines the order in which Aurora

https://docs.amazonaws.cn/en_us/AmazonRDS/latest/AuroraUserGuide/aurora-replicas-adding.html

More than one Aurora Replica can share the same priority, resulting in promotion tiers. If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Backups.html#Aurora.Managing.FaultTolerance>

upvoted 17 times

🗳️ 👤 **dougporto1988** Most Recent 2 years ago

C:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Concepts.AuroraHighAvailability.html>

More than one Aurora Replica can share the same priority, resulting in promotion tiers. If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier.

upvoted 1 times

🗳️ 👤 **f__16** 2 years, 2 months ago

Selected Answer: C

1. Highest priority will be promoted.
2. Largest size will be promoted.
3. Random replica will be promoted.

Since they are the same priority, it will go to (2) Largest size will be promoted.

If all replicas were the same size, a random replica will be promoted.

upvoted 3 times

🗳️ 👤 **sguinales** 1 year, 6 months ago

but he trick here is: The larger instances are reserved for specialized kinds of reporting queries. To make it unlikely for them to be promoted to the primary instance, the following example changes their promotion tier to the lowest priority.

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: C

default tier = -1

More than one Aurora Replica can share the same priority, resulting in promotion tiers. If two or more Aurora Replicas share the same priority, then

Amazon RDS promotes the replica that is largest in size.

upvoted 3 times

🗳️ 👤 **kret** 3 years, 1 month ago

Selected Answer: C

C

"If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier."

upvoted 2 times

🗳️ 👤 **jnassp1** 3 years, 6 months ago

C

If two replicas have the same priority, then the replica that is largest in size gets promoted

upvoted 1 times

🗳️ 👤 **guru_ji** 3 years, 6 months ago

Correct Answer ==>> C

upvoted 1 times

🗳️ 👤 **guru_ji** 3 years, 6 months ago

Correct Answer ==>> C

upvoted 1 times

🗳️ 👤 **LMax** 3 years, 7 months ago

Answer C

upvoted 2 times

🗳️ 👤 **myutran** 3 years, 7 months ago

Ans: C

upvoted 1 times

🗳️ 👤 **Exia** 3 years, 7 months ago

C. If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Concepts.AuroraHighAvailability.html>

upvoted 4 times

🗳️ 👤 **JobinAkaJoe** 3 years, 7 months ago

C is the right choice. Since all replicas have same default priority the largest among them will be considered for failover.

upvoted 2 times

🗳️ 👤 **Billhardy** 3 years, 7 months ago

Since priority isn't assigned, aurora will promote the largest sized replica.

upvoted 1 times

🗳️ 👤 **Ashoks** 3 years, 7 months ago

Ans is C

upvoted 1 times

🗳️ 👤 **Manmohan** 3 years, 7 months ago

Answer: C, If you don't have priority, Aurora will pick the biggest size by default

upvoted 2 times

🗳️ 👤 **Rahu** 3 years, 7 months ago

C is right!

upvoted 4 times

🗳️ 👤 **BillyMadison** 3 years, 8 months ago

<https://aws.amazon.com/blogs/aws/additional-failover-control-for-amazon-aurora/>

A

"In the event of a failover, Amazon RDS will promote the read replica that has the highest priority (the lowest numbered tier). If two or more replicas have the same priority, RDS will promote the one that is the same size as the previous primary instance."

upvoted 2 times

🗳️ 👤 **BillyMadison** 3 years, 7 months ago

Switching to C. AWS has made this question way too hard as their own documentation isn't consistent with their priority levels. I'm going with largest first priority due to <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-ug.pdf>

More than one Aurora Replica can share the same priority, resulting in promotion tiers. If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If

two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier."

upvoted 12 times

A company is running its line of business application on AWS, which uses Amazon RDS for MySQL at the persistent data store. The company wants to minimize downtime when it migrates the database to Amazon Aurora. Which migration method should a Database Specialist use?

- A. Take a snapshot of the RDS for MySQL DB instance and create a new Aurora DB cluster with the option to migrate snapshots.
- B. Make a backup of the RDS for MySQL DB instance using the mysqldump utility, create a new Aurora DB cluster, and restore the backup.
- C. Create an Aurora Replica from the RDS for MySQL DB instance and promote the Aurora DB cluster.
- D. Create a clone of the RDS for MySQL DB instance and promote the Aurora DB cluster.

Suggested Answer: A

Reference:

<https://d1.awsstatic.com/whitepapers/RDS/Migrating%20your%20databases%20to%20Amazon%20Aurora.pdf>


(10)

Community vote distribution

C (100%)

  **chicagomassageseeker** Highly Voted 3 years, 8 months ago

Answer C. Minimal Downtime with Aurora replica. However this is not the fastest.
upvoted 12 times

  **BillyMadison** 3 years, 7 months ago

Agree C.

This blogpost explains why C is the answer since read replicas reduce downtime.

<https://aws.amazon.com/blogs/database/best-practices-for-migrating-rds-for-mysql-databases-to-amazon-aurora/>

upvoted 13 times

  **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: C

C. Create an Aurora Replica from the RDS for MySQL DB instance and promote the Aurora DB cluster.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Migrating.RDSMySQL.Replica.html>

upvoted 1 times

  **f__16** 2 years ago

Selected Answer: C

When migrating from RDS Mysql/Postgres to Aurora,

easiest and most effortless way is to take an (aurora replica and then promote it to standalone).

answer is C.

upvoted 1 times

  **ken_test1234** 2 years, 2 months ago

Selected Answer: C

This article from aws show you that C is exactly the answer <https://aws.amazon.com/getting-started/hands-on/migrate-rdsmysql-to-auroramysql/>

upvoted 1 times

  **Mintwater** 2 years, 2 months ago

C.

Please refer below ---

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Migrating.RDSMySQL.Replica.html>

"

You can migrate from an RDS for MySQL DB instance by first creating an Aurora MySQL read replica of a MySQL DB instance. When the replica lag between the MySQL DB instance and the Aurora MySQL read replica is 0, you can direct your client applications to read from the Aurora read replica and then stop replication to make the Aurora MySQL read replica a standalone Aurora MySQL DB cluster for reading and writing. For details, see

Migrating data from a MySQL DB instance to an Amazon Aurora MySQL DB cluster by using an Aurora read replica.

"

upvoted 1 times

🗳️ 👤 **Anmol_1401** 2 years, 4 months ago

Selected Answer: C

C is the Answer as per cloud guru. Key is minimal downtime

upvoted 2 times

🗳️ 👤 **awsguys** 3 years ago

C => right as minimize downtime.

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: C

MySQL DB instance

-> Create an Aurora Replica (will take time)

-> promote the Aurora DB cluster (Minimal downtime)

upvoted 1 times

🗳️ 👤 **jove** 3 years, 5 months ago

Answer is option C .. Minimal downtime.

upvoted 1 times

🗳️ 👤 **GMartinelli** 3 years, 5 months ago

Selected Answer: C

Option C

upvoted 3 times

🗳️ 👤 **guru_ji** 3 years, 6 months ago

Correct Answer ==>> C

upvoted 2 times

🗳️ 👤 **gelsm** 3 years, 6 months ago

C. Create an Aurora Replica from the RDS for MySQL DB instance and promote the Aurora DB cluster.

This may not be the fastest but the question is "minimize downtime".

upvoted 1 times

🗳️ 👤 **LMax** 3 years, 7 months ago

I go with C

upvoted 1 times

🗳️ 👤 **myutran** 3 years, 7 months ago

Ans: C

upvoted 1 times

🗳️ 👤 **GeeBeeEl** 3 years, 7 months ago

The answer is A and I will prove it --- please follow with me carefully

There is a requirement for minimal downtime (FASTEST) --- I go with A

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Migrating.html#AuroraPostgreSQL.Migrating.RDSPostgreSQL.Import>

According to that link "Be prepared for migration to take a while, roughly several hours per terabyte (TiB) of data." if you go with C. C will thus lead to a longer downtime while A will not

upvoted 2 times

🗳️ 👤 **imatowel** 3 years, 7 months ago

Ok, so what about data written to source RDS after snapshot but before swap to Aurora? You loose data that's unacceptable

upvoted 1 times

🗳️ 👤 **Dip11** 3 years, 7 months ago

You need to understand the difference between minimal downtime and Fastest. Both are not same. Answer is C.

upvoted 1 times

🗳️ 👤 **JobinAkaJoe** 3 years, 7 months ago

C because it facilitates migration to aurora with least down-time.

upvoted 2 times

  **waterh30** 3 years, 7 months ago

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Migrating.html#AuroraPostgreSQL.Migrating.RDSPostgreSQL.Replic>

upvoted 1 times

The Security team for a finance company was notified of an internal security breach that happened 3 weeks ago. A Database Specialist must start producing audit logs out of the production Amazon Aurora PostgreSQL cluster for the Security team to use for monitoring and alerting. The Security team is required to perform real-time alerting and monitoring outside the Aurora DB cluster and wants to have the cluster push encrypted files to the chosen solution.

Which approach will meet these requirements?

- A. Use pg_audit to generate audit logs and send the logs to the Security team.
- B. Use AWS CloudTrail to audit the DB cluster and the Security team will get data from Amazon S3.
- C. Set up database activity streams and connect the data stream from Amazon Kinesis to consumer applications.
- D. Turn on verbose logging and set up a schedule for the logs to be dumped out for the Security team.

Suggested Answer: B

Reference:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-ug.pdf>
(525)

Community vote distribution

C (100%)

  **learnaws**  3 years, 8 months ago

answer is C.

<https://aws.amazon.com/about-aws/whats-new/2019/05/amazon-aurora-with-postgresql-compatibility-supports-database-activity-streams/>
upvoted 15 times

  **BillyC** 3 years, 8 months ago

Yes, real time

upvoted 3 times

  **BillyMadison** 3 years, 7 months ago

Agree C, the link you posted nails it. Anytime the question want "real time alerts or streams", its almost always Kinesis streams.

"Database Activity Streams for Amazon Aurora with PostgreSQL compatibility provides a near real-time data stream of the database activity in your relational database to help you monitor activity. When integrated with third party database activity monitoring tools, Database Activity Streams can monitor and audit database activity to provide safeguards for your database and help meet compliance and regulatory requirements."
upvoted 4 times

  **adelcold**  1 year, 11 months ago

Selected Answer: C

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.Overview.html>
upvoted 1 times

  **ken_test1234** 2 years, 2 months ago

Selected Answer: C

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.Overview.html>
upvoted 1 times

  **SachinGoel** 2 years, 4 months ago

Selected Answer: C

Ans - C

upvoted 1 times

  **Chirantan** 2 years, 10 months ago

answer is C.

<https://aws.amazon.com/about-aws/whats-new/2019/05/amazon-aurora-with-postgresql-compatibility-supports-database-activity-streams/>
Database Activity Streams for Amazon Aurora with PostgreSQL compatibility provides a near real-time data stream of the database activity in your relational database to help you monitor activity. When integrated with third party database activity monitoring tools, Database Activity Streams can monitor and audit database activity to provide safeguards for your database and help meet compliance and regulatory requirements.
upvoted 1 times

🗳️ 👤 **Chirantan** 2 years, 10 months ago

Database Activity Streams for Amazon Aurora with PostgreSQL compatibility provides a near real-time data stream of the database activity in your relational database to help you monitor activity. When integrated with third party database activity monitoring tools, Database Activity Streams can monitor and audit database activity to provide safeguards for your database and help meet compliance and regulatory requirements.

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: C

I chose C over B

x B. Use AWS CloudTrail to audit the DB cluster and the Security team will get data from Amazon S3.

- CloudTrail is automatically enabled for all accounts, and logs events, so this is good for past activities for breach

Amazon Aurora activity is recorded in a CloudTrail event in Event history. You can use the CloudTrail console to view the last 90 days of recorded API activity and events in an AWS Region.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/logging-using-cloudtrail.html>

<https://aws.amazon.com/cloudtrail/faqs/>

C is for future activities

Database Activity Streams can monitor and audit database activity to provide real time safeguards for your database and help meet compliance and regulatory requirements.

<https://aws.amazon.com/about-aws/whats-new/2019/05/amazon-aurora-with-postgresql-compatibility-supports-database-activity-streams/>

DB->Database Activity Streams->Kinesis ->Security team

C. Set up database activity streams and connect the data stream from Amazon Kinesis to consumer applications.

upvoted 2 times

🗳️ 👤 **Mintwater** 2 years, 2 months ago

I like your explanation -- CloudTrail is for the historical data; Kinesis is for the future data.

upvoted 1 times

🗳️ 👤 **jove** 3 years, 5 months ago

Real-time alerting and monitoring > Option C

upvoted 1 times

🗳️ 👤 **Scunningham99** 3 years, 6 months ago

C <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/DBActivityStreams.Overview.html>

upvoted 1 times

🗳️ 👤 **ChauPhan** 3 years, 6 months ago

Only C meets the real-time, A, D is possible but schedule.

upvoted 1 times

🗳️ 👤 **gelsm** 3 years, 7 months ago

C. Set up database activity streams and connect the data stream from Amazon Kinesis to consumer applications.

Aurora Database activity streams provide a near real-time data stream of the database activity for an Aurora DB cluster. Database activity streams require the use of AWS KMS because the activity streams are always encrypted.

upvoted 1 times

🗳️ 👤 **Dip11** 3 years, 7 months ago

C for sure.

upvoted 1 times

🗳️ 👤 **LMax** 3 years, 7 months ago

Between A and C, but would go with C after reading this:

<https://aws.amazon.com/about-aws/whats-new/2019/05/amazon-aurora-with-postgresql-compatibility-supports-database-activity-streams/>

upvoted 2 times

🗳️ 👤 **myutran** 3 years, 7 months ago

Ans: C



upvoted 1 times

🗳️ 👤 **bigaws** 3 years, 7 months ago

I agree with C. I think that Cloudtrail does not support the type of logging that would be required here, it is not the internal database info:
<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Overview.LoggingAndMonitoring.html>
upvoted 1 times

  **JobinAkaJoe** 3 years, 7 months ago

Between A & C I will go with C as the logs should be used for real-time alerting and monitoring.
upvoted 1 times

  **Ashoks** 3 years, 7 months ago

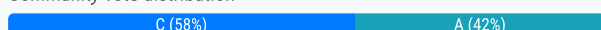
C is the answer
upvoted 1 times

A company is using Amazon RDS for MySQL to redesign its business application. A Database Specialist has noticed that the Development team is restoring their MySQL database multiple times a day when Developers make mistakes in their schema updates. The Developers sometimes need to wait hours for the restores to complete. Multiple team members are working on the project, making it difficult to find the correct restore point for each mistake. Which approach should the Database Specialist take to reduce downtime?

- A. Deploy multiple read replicas and have the team members make changes to separate replica instances
- B. Migrate to Amazon RDS for SQL Server, take a snapshot, and restore from the snapshot
- C. Migrate to Amazon Aurora MySQL and enable the Aurora Backtrack feature
- D. Enable the Amazon RDS for MySQL Backtrack feature

Suggested Answer: A

Community vote distribution



Amogh77 Highly Voted 3 years, 7 months ago

Answer is : A, Backtrack feature is only available with Aurora.

each developer can perform DDL operations for the schema changes on the MySQL read replica once the read replica is in sync with its primary DB instance. Then the developer can promote the read replica and direct the application to use the promoted instance during the development phase. This solution isolates the schema changes done by each developer to their own promoted instance. This also avoids the problem of keeping track of the "correct restore point" that the team faced while using the same DB instance.

upvoted 17 times

saryu 3 years, 7 months ago

Correct

upvoted 2 times

TiredDad 1 year, 11 months ago

Developers would still need to wait hours for the restore to complete, in case of mistakes in their schema updates. But if we try to avoid this, the other option C requires migration to Aurora! and even then "Multiple team members are working on the project, making it difficult to find the correct restore point for each mistake." It is a tough call. <https://serverfault.com/questions/791359/how-do-i-make-an-aws-rds-mysql-read-replica-have-a-different-schema>

this shows how read replica can be used for writes.. with some limitations..

upvoted 1 times

Germaneli 1 year, 8 months ago

That's why C is correct: Migrate to Aurora MySQL, there you can use the BackTrack functionality which is not available in RDS MySQL.

upvoted 1 times

[Removed] Highly Voted 3 years, 8 months ago

I don't see how it can be A. These are read replicas, so by definition you cannot make changes to them. C sounds right to me

upvoted 17 times

chicagomassageseeker 3 years, 8 months ago

yes C is the answer.

upvoted 3 times

BillyMadison 3 years, 8 months ago

C agree

<https://aws.amazon.com/blogs/aws/amazon-aurora-backtrack-turn-back-time/>

upvoted 2 times

Khasta 3 years, 4 months ago

Yes you can write to MySQL readable replicas as mentioned here

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

upvoted 3 times

🗄️ 👤 **Jiang_aws1** 2 years, 7 months ago
if using clone then will be OK but not Replicas - Read only
upvoted 1 times

🗄️ 👤 **MultiAZ** Most Recent 1 year, 4 months ago
Selected Answer: C
Answer is C.
You cannot make changes on a read replica, which makes A invalid.
upvoted 1 times

🗄️ 👤 **Sathish_dbs** 1 year, 7 months ago
the key is 'unable to find the correct restore point', so backtrack won't as you don't know which point it needs backing to as many people making many changes, so first thing is isolate each changes by allocating their own instances
upvoted 1 times

🗄️ 👤 **roymunson** 1 year, 8 months ago
Selected Answer: C
With RDS MySQL it's true to that you can make changes on read replica when setting read_only=0 to make read replica writable, but the use cases are recommended for indexing tables and small maintenance action because other action could break the replication. In addition the option read_only=0 works only for engines like InnoDB and with MyISAM u break the replication. In the question is nothing written about the engine type and that's why C is the right answer in my opinion.
upvoted 1 times

🗄️ 👤 **roymunson** 1 year, 8 months ago
Mmmhh now I saw the exact same question in an Udemy Exam Course and the authors are saying it is indeed Answer A with the following explanation:

"For the given use-case, each developer can perform DDL operations for the schema changes on the MySQL read replica once the read replica is in sync with its primary DB instance. Then the developer can promote the read replica and direct the application to use the promoted instance during the development phase. This solution isolates the schema changes done by each developer to their own promoted instance. This also avoids the problem of keeping track of the "correct restore point" that the team faced while using the same DB instance.

If you need to make changes to the MySQL or MariaDB read replica, you must set the read_only parameter to 0 in the DB parameter group for the read replica. You can then perform all needed DDL operations, such as creating indexes, on the read replica. Actions taken on the read replica don't affect the performance of the primary DB instance."
upvoted 1 times

🗄️ 👤 **aqiao** 1 year, 8 months ago
Selected Answer: C
Although you can enable updates by setting the read_only parameter to 0 in the DB parameter group for the read replica, we recommend that you don't do so problems if the read replica becomes incompatible with the source DB instance
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_MySQL.Replication.ReadReplicas.html#USER_MySQL.Replication.ReadReplicas.DelayRe
Beside, how to find the restore point with replica after mistake
upvoted 2 times

🗄️ 👤 **IhorK** 1 year, 9 months ago
Selected Answer: A
A is the answer.
"...have the team members make changes to separate replica instances..." - we can make MySQL read replica writable. This will destroy replication, but for our purposes it is not critical.
C - "Aurora Backtrack feature" doesn't suite because of "Multiple team members are working on the project, making it difficult to find the correct restore point for each mistake."
upvoted 1 times

🗄️ 👤 **Monknil** 1 year, 10 months ago
Tricky one, is it C or A ? What is the Final answer some say A some C. Both seem plausible, but which is the better option
upvoted 2 times

🗄️ 👤 **mraronsimon** 1 year, 10 months ago
Selected Answer: A
A (correct)
B - incorrect - Migration to SQL server needs a high impact on the application code. Restoring from a snapshot is much slower than using a prepared read replica.

C - incorrect - I think there was a good reason why the company use RDS instead of Aurora. But if you migrate to Aurora better to use the Clone feature instead of Backtrack in that case.

D - incorrect - Backtrack is an Aurora (MySQL only) feature, not for RDS.

upvoted 1 times

🗨️ 👤 **megramlak** 2 years ago

Answer is A because this involves minimum effort

upvoted 1 times

🗨️ 👤 **Kodoma** 2 years ago

Selected Answer: A

As stated by others, PITR is not a good option since it is "difficult to find the correct restore point for each mistake", and PITR is also a feature only available on Aurora

upvoted 1 times

🗨️ 👤 **MrAliMohsan** 2 years ago

Selected Answer: A

The answer is A: Since MySQL Read replicas can be made writeable as per below documentation link. Previously, I was also thinking of C as a correct answer team members are involved in the project, which makes it difficult to identify the right restoration point for each error"

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html#:~:text=Yes.%C2%A0You%20can%20enable%20the%20MySQL%20or%20

upvoted 2 times

🗨️ 👤 **ken_test1234** 2 years, 2 months ago

Selected Answer: C

How is it possible to update or modify a read replica ?

upvoted 2 times

🗨️ 👤 **Theprince_Sree** 2 years, 3 months ago

Selected Answer: C

Answer is C: Migrate to Aurora & use back track feature.

upvoted 2 times

🗨️ 👤 **Theprince_Sree** 2 years, 3 months ago

Answer is C: Solution is to migrate the RDS to Aurora and enable backtrack. Folks who are saying A needs to explain on how they can apply changes on Read replicas and rollback them!!??

upvoted 1 times

🗨️ 👤 **im_not_robot** 2 years, 3 months ago

Selected Answer: C

A is incorrect because 2 reasons: the first reason is read replica is read-only. The second reason is you need to modify the app to use different read replica instance endpoint for each team. If they have only one business app connects to a database, it is nightmare to connect to multiple read replica for each team

B is incorrect because it is more expensive and can not solve the problem

C is correct at backtrack is in-place restore

D is incorrect because rds mysql doesn't support backtrack

sorry for my english

upvoted 2 times

🗨️ 👤 **awsexams** 2 years, 6 months ago

Selected Answer: A

I will go with A as well, mainly because of the wording on "correct restore point"

upvoted 3 times

A media company is using Amazon RDS for PostgreSQL to store user data. The RDS DB instance currently has a publicly accessible setting enabled and is hosted in a public subnet. Following a recent AWS Well-Architected Framework review, a Database Specialist was given new security requirements.

- ⇒ Only certain on-premises corporate network IPs should connect to the DB instance.
- ⇒ Connectivity is allowed from the corporate network only.

Which combination of steps does the Database Specialist need to take to meet these new requirements? (Choose three.)

- A. Modify the pg_hba.conf file. Add the required corporate network IPs and remove the unwanted IPs.
- B. Modify the associated security group. Add the required corporate network IPs and remove the unwanted IPs.
- C. Move the DB instance to a private subnet using AWS DMS.
- D. Enable VPC peering between the application host running on the corporate network and the VPC associated with the DB instance.
- E. Disable the publicly accessible setting.
- F. Connect to the DB instance using private IPs and a VPN.

Suggested Answer: DEF

Community vote distribution

BEF (100%)

🗳️ **Ebi** Highly Voted 3 years, 7 months ago

Answer is BEF

Database does not need to be in private subnet (there is no requirement in the question) disabling public accessibility will remove public IP address associated from the instance.

upvoted 14 times

🗳️ **ChauPhan** Highly Voted 3 years, 7 months ago

B E F are the answers.

D is not correct because there is on-premise network, VPC peering is for AWS VPC - AWS VPC

C is not correct. DMS is using for DB migration, not subnet modification

upvoted 8 times

🗳️ **guru_ji** 3 years, 6 months ago

Correct Answer ==> B,E,F

upvoted 1 times

🗳️ **IhorK** Most Recent 1 year, 9 months ago

Selected Answer: BEF

<https://kloudle.com/academy/how-to-restrict-access-to-your-publicly-accessible-rds-instance/>

After publicly accessible setting disable no need to move the DB instance to a private subnet.

We do not change the conf file settings, instead we change the security group where we configure the IP addresses from which access is required.

VPC peering only inside AWS.

You need to select the third item, nothing but "Connect to the DB instance using private IPs and a VPN" is suitable.

upvoted 1 times

🗳️ **ankurlibra** 2 years, 3 months ago

BEF for sure

upvoted 1 times

🗳️ **novice_expert** 3 years ago

Selected Answer: BEF

x A. RDS you don't edit config files directly

B. Modify the security group. Add the required corporate network IPs and remove the unwanted IPs

x C. subnet change by DMS?

x D. VPC peering is within AWS only

E. disable publicly accessible

F. .Connect to the DB instance using private IPs and a VPN.

upvoted 1 times

🗳️ 👤 **kush_sumit** 2 years, 10 months ago

You cant ssh directly into rds how would you connect using private IP's?

upvoted 1 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: BEF

1) Security Groups HAS to be done to restrict DB access to specific IPS

2) Public accessibility has to be removed

3) Corp to AWS VPN has to be enabled to secure traffic

upvoted 2 times

🗳️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: BEF

agree with others

upvoted 2 times

🗳️ 👤 **awsmonster** 3 years, 4 months ago

F is incorrect. RDS uses endpoint, not IP address.

I vote for BCE

upvoted 3 times

🗳️ 👤 **GMartinelli** 3 years, 5 months ago

Selected Answer: BEF

B, E & F

upvoted 3 times

🗳️ 👤 **manan728** 3 years, 7 months ago

B,C and E are correct. You need to migrate the database to private subnet before you can disable the publicly accessible setting in the console.

upvoted 2 times

🗳️ 👤 **Windy** 3 years, 7 months ago

BEF for me

upvoted 3 times

🗳️ 👤 **myutran** 3 years, 7 months ago

Ans: BEF

upvoted 3 times

🗳️ 👤 **JobinAkaJoe** 3 years, 7 months ago

I will go with BEF.

Ideally db should be moved to private subnet. But using DMS for that makes C a wrong choice

upvoted 4 times

🗳️ 👤 **kilkar** 3 years, 7 months ago

BDF

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-connectivity-instance-subnet-vpc/>

https://docs.aws.amazon.com/vpn/latest/s2svpn/VPC_VPN.html

upvoted 1 times

🗳️ 👤 **ChauPhan** 3 years, 7 months ago

VPC peering is between AWS VPCs, not between on-premise network and AWS VPC

upvoted 1 times

🗳️ 👤 **[Removed]** 3 years, 7 months ago

F - Connect to DB instance using VPC peering - At which point in this question does it mention the need to connect to other VPC's? Best practices in AWS doco state that (see the Note) - -

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.WorkingWithRDSInstanceinaVPC.html#USER_VPC.Hiding

upvoted 1 times

🗳️ 👤 **Ashoks** 3 years, 7 months ago

B,C,E,F are probable answers.

C - database needs to moved private to public subnet, however, migration can be done through snapshot instead dms.

So BEF would be the answers

upvoted 2 times

  **Manmohan** 3 years, 7 months ago

BEF for me

upvoted 1 times

A company is about to launch a new product, and test databases must be re-created from production data. The company runs its production databases on an Amazon Aurora MySQL DB cluster. A Database Specialist needs to deploy a solution to create these test databases as quickly as possible with the least amount of administrative effort. What should the Database Specialist do to meet these requirements?

- A. Restore a snapshot from the production cluster into test clusters
- B. Create logical dumps of the production cluster and restore them into new test clusters
- C. Use database cloning to create clones of the production cluster
- D. Add an additional read replica to the production cluster and use that node for testing

Suggested Answer: D

Community vote distribution

C (100%)


  **BillyMadison** Highly Voted 3 years, 7 months ago

C.

<https://aws.amazon.com/getting-started/hands-on/aurora-cloning-backtracking/>

"Cloning an Aurora cluster is extremely useful if you want to assess the impact of changes to your database, or if you need to perform workload-intensive operations—such as exporting data or running analytical queries, or simply if you want to use a copy of your production database in a development or testing environment. You can make multiple clones of your Aurora DB cluster. You can even create additional clones from other clones, with the constraint that the clone databases must be created in the same region as the source databases.

upvoted 17 times

  **Ihork** Most Recent 1 year, 9 months ago

Selected Answer: C

When we do cloning, there is no physical overwriting of data to a new location, only meta-data is copied. Therefore, the cloning process is quite fast. However, further changes are isolated from each other in source and clone.

upvoted 1 times

  **ken_test1234** 2 years, 2 months ago

Selected Answer: C

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Clone.html>

upvoted 1 times

  **f__16** 2 years, 2 months ago

Selected Answer: C

Cloning is the best choice. It creates a new db cluster with the same data.

upvoted 1 times

  **Jiang_aws1** 2 years, 7 months ago

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Clone.html>

Creating a clone is faster and more space-efficient than physically copying the data using other techniques, such as restoring a snapshot.

upvoted 2 times

  **novice_expert** 3 years ago

Selected Answer: C

C. Use database cloning to create clones of the production cluster

Cloning is best choice with "copy-on-write" protocol, database becomes available in a few mins.

<https://aws.amazon.com/getting-started/hands-on/aurora-cloning-backtracking/>

"Cloning an Aurora cluster is extremely useful if you want to assess the impact of changes to your database, or if you need to perform workload-intensive operations—such as exporting data or running analytical queries, or simply if you want to use a copy of your production database in a development or testing environment. You can make multiple clones of your Aurora DB cluster. You can even create additional clones from other clones, with the constraint that the clone databases must be created in the same region as the source databases.

upvoted 2 times

🗳️ 👤 **AriraAWS** 3 years, 3 months ago

Selected Answer: C

Cloning is best choice with "copy-on-write" protocol, database becomes available in a few mins.

upvoted 2 times

🗳️ 👤 **LMax** 3 years, 6 months ago

Answer C

upvoted 2 times

🗳️ 👤 **myutran** 3 years, 6 months ago

Ans: C

upvoted 1 times

🗳️ 👤 **JobinAkaJoe** 3 years, 6 months ago

Answer is C

Aurora copy-on-write clones are best suited for this requirement

upvoted 1 times

🗳️ 👤 **BillyC** 3 years, 7 months ago

Yes, C

upvoted 2 times

🗳️ 👤 **jnassp1** 3 years, 7 months ago

C Database Clone is quickest and the right approach. Takes secs..

upvoted 3 times

🗳️ 👤 **Mickysingh** 3 years, 7 months ago

C is correct as we have least admin effort and quck

upvoted 3 times

🗳️ 👤 **BillyC** 3 years, 7 months ago

D here!

upvoted 1 times

🗳️ 👤 **[Removed]** 3 years, 7 months ago

D is neither least administrative nor will it meet the requirement. A is simplest for me

upvoted 2 times

🗳️ 👤 **Ebi** 3 years, 7 months ago

Restoring from snapshot is never quick.

Answer is C

upvoted 5 times

A company with branch offices in Portland, New York, and Singapore has a three-tier web application that leverages a shared database. The database runs on

Amazon RDS for MySQL and is hosted in the us-west-2 Region. The application has a distributed front end deployed in the us-west-2, ap-southeast-1, and us-east-2 Regions.

This front end is used as a dashboard for Sales Managers in each branch office to see current sales statistics. There are complaints that the dashboard performs more slowly in the Singapore location than it does in Portland or New York. A solution is needed to provide consistent performance for all users in each location.



Which set of actions will meet these requirements?

- A. Take a snapshot of the instance in the us-west-2 Region. Create a new instance from the snapshot in the ap-southeast-1 Region. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- B. Create an RDS read replica in the ap-southeast-1 Region from the primary RDS DB instance in the us-west-2 Region. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- C. Create a new RDS instance in the ap-southeast-1 Region. Use AWS DMS and change data capture (CDC) to update the new instance in the ap-southeast-1 Region. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.
- D. Create an RDS read replica in the us-west-2 Region where the primary instance resides. Create a read replica in the ap-southeast-1 Region from the read replica located on the us-west-2 Region. Reconfigure the ap-southeast-1 front-end dashboard to access this instance.

Suggested Answer: A

Community vote distribution

B (100%)

  **BillyMadison** Highly Voted 3 years, 7 months ago

Leaning to B.

<https://aws.amazon.com/rds/features/read-replicas/>

"Amazon RDS Read Replicas provide enhanced performance and durability for RDS database (DB) instances. They make it easy to elastically scale out beyond the capacity constraints of a single DB instance for read-heavy database workloads. You can create one or more replicas of a given source DB Instance and serve high-volume application read traffic from multiple copies of your data, thereby increasing aggregate read throughput. "

upvoted 20 times

  **khchan123** 3 years ago

Agree it's B.

upvoted 1 times

  **[Removed]** Highly Voted 3 years, 7 months ago

If they go with A, how do they keep the databases in sync?

B for me

upvoted 6 times

  **NishithShah** Most Recent 1 year, 5 months ago

Selected Answer: B



B is correct.

upvoted 1 times

  **Sathish_dbs** 1 year, 7 months ago



apparently I need to know geography as well that Singapore is AP region

upvoted 1 times

  **jitesh_k** 1 year, 6 months ago

They assume as much.

upvoted 1 times

  **Germaneli** 1 year, 8 months ago

Selected Answer: B

From the plausible options,

A - this would make it a non-shared database, destroying the architecture.

B - iff the dashboard is only for reading then the read replica in SGP might do.

upvoted 1 times

🗳️ 👤 **ken_test1234** 2 years, 2 months ago

Selected Answer: B

Because the read replica is designed for this purpose

upvoted 1 times

🗳️ 👤 **im_not_robot** 2 years, 3 months ago

Selected Answer: B

A is incorrect in no data synchronization between 2 databases

B is correct since the read replica is same region with the frontend

C is incorrect because extra cost of DMS instance. Although it is feasible option.

D is incorrect because of extra cost of new replica on primary region and replication lag.

upvoted 3 times

🗳️ 👤 **examineme** 2 years, 5 months ago

Selected Answer: B

B is the answer

upvoted 1 times

🗳️ 👤 **sachin** 2 years, 10 months ago

You can create read replication from a already existing replica, but the replication lag will be high. So D is operationally possible but replication lag will be high.

So correct answer is B.

upvoted 1 times

🗳️ 👤 **ryuhei** 2 years, 11 months ago

Selected Answer: B

Answer:B

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: B

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.XRgn.html

You can create replicas through Region same way as in same region

Create an RDS read replica in the ap-southeast-1 Region from the primary RDS DB instance in the us-west-2 Region. Reconfigure the ap-southeast-1 front- end dashboard to access this instance.

upvoted 1 times

🗳️ 👤 **AriraAWS** 3 years, 3 months ago

Selected Answer: B

B is the right answer!!

upvoted 2 times

🗳️ 👤 **ChauPhan** 3 years, 6 months ago

B is correct. You can create replicas through Region

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.XRgn.html

upvoted 1 times

🗳️ 👤 **Dip11** 3 years, 6 months ago

I will go with B

upvoted 1 times

🗳️ 👤 **LMax** 3 years, 6 months ago

Would go with B

upvoted 2 times

🗳️ 👤 **myutran** 3 years, 6 months ago

Ans: B

upvoted 1 times

🗳️ 👤 **Bassel** 3 years, 6 months ago

Not sure why you all excluded C. It is definitely a good solution to keep two databases in sync using DMS with CDC.

upvoted 3 times

A company wants to migrate its existing on-premises Oracle database to Amazon Aurora PostgreSQL. The migration must be completed with minimal downtime using AWS DMS. A Database Specialist must validate that the data was migrated accurately from the source to the target before the cutover. The migration must have minimal impact on the performance of the source database.

Which approach will MOST effectively meet these requirements?

- A. Use the AWS Schema Conversion Tool (AWS SCT) to convert source Oracle database schemas to the target Aurora DB cluster. Verify the datatype of the columns.
- B. Use the table metrics of the AWS DMS task created for migrating the data to verify the statistics for the tables being migrated and to verify that the data definition language (DDL) statements are completed.
- C. Enable the AWS Schema Conversion Tool (AWS SCT) premigration validation and review the premigration checklist to make sure there are no issues with the conversion.
- D. Enable AWS DMS data validation on the task so the AWS DMS task compares the source and target records, and reports any mismatches.

Suggested Answer: D

Reference:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Validating.html

Data validation

Tagging resources

Security

Limits

▶ Troubleshooting and diagnostic support

▶ Migrating large data stores with Snowball Edge

▶ Reference

Release notes

Document history

AWS glossary


Community vote distribution

D (100%)

• Limitations

AWS DMS provides support for data validation to ensure that your data was migrated accurately from the source to the target. If enabled, validation begins immediately after a full load is performed for a table. Validation compares the incremental changes for a CDC-enabled task as they occur.

During data validation, AWS DMS compares each row in the source with its corresponding row at the target, verifies the rows contain the same data, and reports any mismatches. To accomplish this AWS DMS issues appropriate queries to retrieve the data. Note that these queries will consume additional resources at the source and target as well as additional network resources.

 **grekh001** Highly Voted 3 years, 5 months ago

"To ensure that your data was migrated accurately from the source to the target, we highly recommend that you use data validation."

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html

Answer is D.

upvoted 9 times

 **Hisayuki** Most Recent 1 year, 4 months ago

Selected Answer: D

DMS supports the validation between source and target data. But SCT does not support it.

upvoted 1 times

 **sachin** 2 years, 11 months ago

table metrics of the AWS DMS can be verified manually for DMS task. There is no mechanism which provided an automated for reading these metrics and providing confirmation if everything went well. If you are working 100+ tables then DMS Data Validation is the only option.

D is best fit

upvoted 2 times

 **novice_expert** 3 years ago

Selected Answer: D

D would run select queries on source & target to compare rows, so some load

B is also good candidate, rather better but it needed "The data validation option in the DMS task has to be activated before DMS "

You can find individual table metrics on the Table statistics tab for each individual task. These metrics include these numbers:

Rows loaded during the full load.

Inserts, updates, and deletes since the task started.

DDL operations since the task started.

B. Use the table metrics of the AWS DMS task created for migrating the data to verify the statistics for the tables being migrated and to verify that the data definition language (DDL) statements are completed.

D. Enable AWS DMS data validation on the task so the AWS DMS task compares the source and target records, and reports any mismatches.

upvoted 2 times

🗳️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: D

D for data validation

upvoted 3 times

🗳️ 👤 **SMAZ** 3 years, 5 months ago

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Validating.html

During data validation, AWS DMS compares each row in the source with its corresponding row at the target, verifies the rows contain the same data, and reports any mismatches. To accomplish this AWS DMS issues appropriate queries to retrieve the data. Note that these queries will consume additional resources at the source and target as well as additional network resources.

So Answer should be B

upvoted 1 times

🗳️ 👤 **awsmonster** 3 years, 4 months ago

Answer should be D:

The data validation option in the DMS task has to be activated before DMS performs what SMAZ has written. B does not mention anything about enabling it.

upvoted 3 times

🗳️ 👤 **SMAZ** 3 years, 5 months ago

'The migration should have a negligible effect on the source database's performance.'

I believe answer should be 'B'

Table metrics

You can find individual table metrics on the Table statistics tab for each individual task. These metrics include these numbers:

Rows loaded during the full load.

Inserts, updates, and deletes since the task started.

DDL operations since the task started.

upvoted 2 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

No datavalidation can ONLY be done using Task Validation and it has to be enabled before DMS tasks start and after migration is finished. Its a very common question. D is the correct choice

upvoted 1 times

🗳️ 👤 **2025flakyt** 3 years, 5 months ago

D is the correct answer

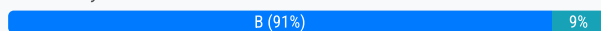
upvoted 3 times

A company is planning to close for several days. A Database Specialist needs to stop all applications along with the DB instances to ensure employees do not have access to the systems during this time. All databases are running on Amazon RDS for MySQL. The Database Specialist wrote and ran a script to stop all the DB instances. When reviewing the logs, the Database Specialist found that Amazon RDS DB instances with read replicas did not stop. How should the Database Specialist edit the script to fix this issue?

- A. Stop the source instances before stopping their read replicas
- B. Delete each read replica before stopping its corresponding source instance
- C. Stop the read replicas before stopping their source instances
- D. Use the AWS CLI to stop each read replica and source instance at the same time

Suggested Answer: D

Community vote distribution



awscamus Highly Voted 3 years, 8 months ago

B. Because You can't stop a DB instance that has a read replica, or that is a read replica. https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_StopInstance.html
upvoted 19 times

aviathor 2 years ago

Where does it say that you need to stop the source instance before read replicas?
upvoted 1 times

aviathor 2 years ago

Found it
upvoted 1 times

adelcold Most Recent 1 year, 11 months ago

Selected Answer: B

You can't stop a DB instance that has a read replica, or that is a read replica.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_StopInstance.html
upvoted 1 times

ken_test1234 2 years, 2 months ago

Selected Answer: B

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_StopInstance.html
upvoted 1 times

aviathor 2 years ago

Where does it say that you need to stop the source instance before read replicas?
upvoted 1 times

aviathor 2 years ago

Found it
upvoted 1 times

sayed 2 years, 7 months ago

Selected Answer: B

B because

- Can stop an RDS instance only if it does not have a replica
 - Cannot stop an RDS replica
- upvoted 2 times

awsjij 2 years, 7 months ago

Selected Answer: B

You can't stop a DB instance that has a read replica, or that is a read replica

upvoted 1 times

🗨️ **awsguys** 3 years ago

B. as terminated all database instances

upvoted 1 times

🗨️ **novice_expert** 3 years ago

Selected Answer: D

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/SQLServer.Concepts.General.SSL.Using.html>

D. addressing the port where the RDS DB instance is listening for encrypted connections (communications connection failure => port)

upvoted 1 times

🗨️ **novice_expert** 3 years ago

sorry it is for next question.

For this one answer is B

upvoted 2 times

🗨️ **novice_expert** 3 years ago

Selected Answer: B

B. Delete each read replica before stopping its corresponding source instance

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_StopInstance.html

Limitations

The following are some limitations to stopping and starting a DB instance:

You can't stop a DB instance that has a read replica, or that is a read replica.

You can't stop an Amazon RDS for SQL Server DB instance in a Multi-AZ configuration.

You can't modify a stopped DB instance.

You can't delete an option group that is associated with a stopped DB instance.

You can't delete a DB parameter group that is associated with a stopped DB instance.

In a Multi-AZ configuration, the primary and secondary Availability Zones might be switched after you start the DB instance.

upvoted 3 times

🗨️ **kret** 3 years, 1 month ago

Selected Answer: B

"You can't stop a DB instance that has a read replica, or that is a read replica."

upvoted 2 times

🗨️ **ChauPhan** 3 years, 6 months ago

B is correct

Limitations

The following are some limitations to stopping and starting a DB instance:

You can't stop a DB instance that has a read replica, or that is a read repl

upvoted 1 times

🗨️ **LMax** 3 years, 6 months ago

Answer B

upvoted 1 times

🗨️ **myutran** 3 years, 6 months ago

Ans: B

upvoted 1 times

🗨️ **JobinAkaJoe** 3 years, 7 months ago

B is correct

upvoted 1 times

🗨️ 👤 **Billhardy** 3 years, 7 months ago

Answer is B. The read replicas need to be deleted in order to stop the primary.

upvoted 1 times

🗨️ 👤 **Ashoks** 3 years, 7 months ago

Ans is B

upvoted 1 times

🗨️ 👤 **BillyC** 3 years, 7 months ago

Ans B is correct

upvoted 2 times

🗨️ 👤 **BillyMadison** 3 years, 8 months ago

I think B.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_StopInstance.html

"The following are some limitations to stopping and starting a DB instance:

You can't stop a DB instance that has a read replica, or that is a read replica."

So if you cant stop a db with a read replica, you have to delete the read replica first to then stop it???

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_MySQL.Replication.ReadReplicas.html#USER_MySQL.Replication.ReadReplicas.StartStop

upvoted 3 times

A global digital advertising company captures browsing metadata to contextually display relevant images, pages, and links to targeted users. A single page load can generate multiple events that need to be stored individually. The maximum size of an event is 200 KB and the average size is 10 KB. Each page load must query the user's browsing history to provide targeting recommendations. The advertising company expects over 1 billion page visits per day from users in the United States, Europe, Hong Kong, and India. The structure of the metadata varies depending on the event. Additionally, the browsing metadata must be written and read with very low latency to ensure a good viewing experience for the users. Which database solution meets these requirements?

- A. Amazon DocumentDB
- B. Amazon RDS Multi-AZ deployment
- C. Amazon DynamoDB global table
- D. Amazon Aurora Global Database

Suggested Answer: C

Reference:

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

novice_expert Highly Voted 3 years ago

Selected Answer: C

- information structure differs according to the event => SchemaLess

Amazon DynamoDB global table

upvoted 8 times

GMartinelli Highly Voted 3 years, 5 months ago

Selected Answer: C

Option C

upvoted 7 times

Pranava_GCP Most Recent 1 year, 8 months ago

Selected Answer: C

C. Amazon DynamoDB global table

"DynamoDB global tables are ideal for massively scaled applications with globally dispersed users. In such an environment, users expect very fast application performance. Global tables provide automatic multi-active replication to AWS Regions worldwide. They enable you to deliver low-latency data access to your users no matter where they are located."

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

upvoted 3 times

🗨️ 👤 **IhorK** 1 year, 9 months ago

Selected Answer: C

- The structure of the metadata varies
- written and read with VERY LOW latency
- average size is 10 KB
- over 1 billion page visits per day
- users in the United States, Europe, Hong Kong, and India

-> C. Amazon DynamoDB global table

upvoted 3 times

🗨️ 👤 **yahoos** 3 years, 3 months ago

C

DynamoDB global table

upvoted 1 times

🗨️ 👤 **jove** 3 years, 5 months ago

Go with option C

upvoted 2 times

A Database Specialist modified an existing parameter group currently associated with a production Amazon RDS for SQL Server Multi-AZ DB instance. The change is associated with a static parameter type, which controls the number of user connections allowed on the most critical RDS SQL Server DB instance for the company. This change has been approved for a specific maintenance window to help minimize the impact on users.

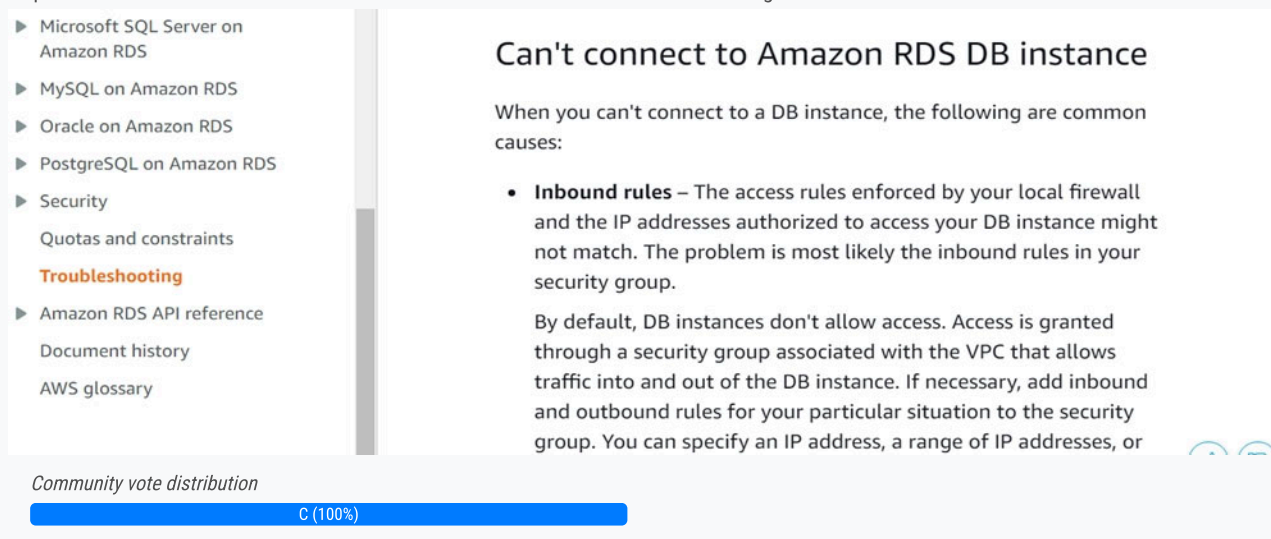
How should the Database Specialist apply the parameter group change for the DB instance?

- A. Select the option to apply the change immediately
- B. Allow the preconfigured RDS maintenance window for the given DB instance to control when the change is applied
- C. Apply the change manually by rebooting the DB instance during the approved maintenance window
- D. Reboot the secondary Multi-AZ DB instance

Suggested Answer: D

Reference:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Troubleshooting.html



Can't connect to Amazon RDS DB instance

When you can't connect to a DB instance, the following are common causes:

- Inbound rules** – The access rules enforced by your local firewall and the IP addresses authorized to access your DB instance might not match. The problem is most likely the inbound rules in your security group.

By default, DB instances don't allow access. Access is granted through a security group associated with the VPC that allows traffic into and out of the DB instance. If necessary, add inbound and outbound rules for your particular situation to the security group. You can specify an IP address, a range of IP addresses, or

Community vote distribution

C (100%)

 **RotterDam** Highly Voted 3 years, 2 months ago


Selected Answer: C

C - static parameters require a manual reboot. You can either do it immediately or during the maintenance window
upvoted 7 times

 **hgelli** Most Recent 1 year, 8 months ago

Selected Answer: C

Static parameters require a reboot. So the best option is rebooting it during the approved maintenance window.
upvoted 1 times

 **IhorK** 1 year, 9 months ago

Selected Answer: C

Static and dynamic DB parameters could be for instance or cluster.
Static parameters change takes effect ONLY after you manually reboot the associated DB clusters or DB instances according to documentation.
<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/parameter-groups-overview.html>
upvoted 1 times

 **navkumin** 1 year, 11 months ago

Selected Answer: C

C - static parameters require a manual reboot. For minimal impact, wait until the maintenance window.
upvoted 1 times

 **Mintwater** 2 years, 2 months ago

C -
<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/parameter-groups-overview.html#parameter-groups-overview.db-instance>

When you change a static parameter and save the DB parameter group, the parameter change takes effect after you manually reboot the associated DB instances. For static parameters, the console always uses pending-reboot for the ApplyMethod

upvoted 1 times

🗲️ 👤 **sayed** 2 years, 7 months ago

Selected Answer: C

C

changing static parameters in a Parameter Group requires a manual reboot, RDS will not reboot the instance during maintenance window.

upvoted 1 times

🗲️ 👤 **novice_expert** 3 years ago

Selected Answer: C

C. Apply the change manually by rebooting the DB instance during the approved maintenance window

upvoted 3 times

🗲️ 👤 **keitahigaki** 3 years, 6 months ago

Ans C

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithParamGroups.html#USER_WorkingWithParamGroups.Modifying

upvoted 2 times

🗲️ 👤 **toppic26** 3 years, 6 months ago

Answer is C

upvoted 1 times

A Database Specialist is designing a new database infrastructure for a ride hailing application. The application data includes a ride tracking system that stores GPS coordinates for all rides. Real-time statistics and metadata lookups must be performed with high throughput and microsecond latency. The database should be fault tolerant with minimal operational overhead and development effort. Which solution meets these requirements in the MOST efficient way?

- A. Use Amazon RDS for MySQL as the database and use Amazon ElastiCache
- B. Use Amazon DynamoDB as the database and use DynamoDB Accelerator
- C. Use Amazon Aurora MySQL as the database and use Aurora's buffer cache
- D. Use Amazon DynamoDB as the database and use Amazon API Gateway

Suggested Answer: D

Reference:

<https://aws.amazon.com/solutions/case-studies/lyft/>

Community vote distribution

B (100%)

 **learnaws**  3 years, 8 months ago

B. keyword here is microsecond. DAX has that
upvoted 15 times

 **BillyMadison** 3 years, 8 months ago

Agree, microsecond = DynamoDB Dax

[https://aws.amazon.com/dynamodb/dax/#:~:text=Amazon%20DynamoDB%20Accelerator%20\(DAX\)%20is,millions%20of%20requests%20per%20second.](https://aws.amazon.com/dynamodb/dax/#:~:text=Amazon%20DynamoDB%20Accelerator%20(DAX)%20is,millions%20of%20requests%20per%20second.)

"Amazon DynamoDB Accelerator (DAX) is a fully managed, highly available, in-memory cache for DynamoDB that delivers up to a 10x performance improvement – from milliseconds to microseconds – even at millions of requests per second. "

upvoted 6 times

 **hgelli**  1 year, 8 months ago

Selected Answer: B

Microsecond latency (i.e., Elastic Cache or DX).

Little operational overhead, DynamoDB.

upvoted 4 times

 **SachinGoel** 2 years, 4 months ago

Selected Answer: B

Unstructured + microsecond - Dynamo + DAX

upvoted 2 times

 **sachin** 2 years, 10 months ago

Since the data lookup should be consistent and real time ,so eventual consistency will not work. DAX will not help . I think D is correct

upvoted 2 times

 **Dantas** 2 years, 11 months ago

Selected Answer: B

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


upvoted 1 times

 **novice_expert** 3 years ago

Selected Answer: B

microsecond => DAX

upvoted 1 times

 **GMartinelli** 3 years, 5 months ago

Selected Answer: B

Option B

upvoted 1 times

🗨️ 👤 **GMartinelli** 3 years, 5 months ago

Selected Answer: B

Option B

upvoted 1 times

🗨️ 👤 **LMax** 3 years, 6 months ago

Agree with B due to latency requirement that can be full field only with DAX

upvoted 1 times

🗨️ 👤 **myutran** 3 years, 7 months ago

Ans: B

upvoted 1 times

🗨️ 👤 **JobinAkaJoe** 3 years, 7 months ago

B - NoSQL is ideal here. DynamoDB+DAX provides microseconds latency

upvoted 4 times

🗨️ 👤 **Billhardy** 3 years, 7 months ago

I will go with B. DAX facilitates the requirement stated

upvoted 2 times

🗨️ 👤 **Ashoks** 3 years, 7 months ago

Ans B. Low latency = DynamoDB, Micro sec = DAX

upvoted 1 times

🗨️ 👤 **halol** 3 years, 7 months ago

Answer is B ,

upvoted 1 times

🗨️ 👤 **Ebi** 3 years, 7 months ago

Answer is B,

Microsecond latency is always NoSQL

API Gateway has nothing to do with requirement of this question.

upvoted 3 times

🗨️ 👤 **Mickysingh** 3 years, 8 months ago

D is correct as document says API gateway to use for minimize developer efforts

upvoted 2 times

🗨️ 👤 **[Removed]** 3 years, 8 months ago

I looked at the link, and watched the video. At no point was API gateway mentioned, Dynamo was.

I will go with B

upvoted 2 times

A company is using an Amazon Aurora PostgreSQL DB cluster with an xlarge primary instance master and two large Aurora Replicas for high availability and read-only workload scaling. A failover event occurs and application performance is poor for several minutes. During this time, application servers in all Availability

Zones are healthy and responding normally.

What should the company do to eliminate this application performance issue?

- A. Configure both of the Aurora Replicas to the same instance class as the primary DB instance. Enable cache coherence on the DB cluster, set the primary DB instance failover priority to tier-0, and assign a failover priority of tier-1 to the replicas.
- B. Deploy an AWS Lambda function that calls the DescribeDBInstances action to establish which instance has failed, and then use the PromoteReadReplica operation to promote one Aurora Replica to be the primary DB instance. Configure an Amazon RDS event subscription to send a notification to an Amazon SNS topic to which the Lambda function is subscribed.
- C. Configure one Aurora Replica to have the same instance class as the primary DB instance. Implement Aurora PostgreSQL DB cluster cache management. Set the failover priority to tier-0 for the primary DB instance and one replica with the same instance class. Set the failover priority to tier-1 for the other replicas.
- D. Configure both Aurora Replicas to have the same instance class as the primary DB instance. Implement Aurora PostgreSQL DB cluster cache management. Set the failover priority to tier-0 for the primary DB instance and to tier-1 for the replicas.

Suggested Answer: D

Community vote distribution

C (100%)

  **RBSK**  2 years, 4 months ago

Selected Answer: C

Tier-0 is a requirement for CCM to work. Also having more than 1 read replica with Tier-0 also disables CCM - https://docs.amazonaws.cn/en_us/AmazonRDS/latest/AuroraUserGuide/aurora_ccm_status.html



This clearly eliminates Option-D
upvoted 5 times

  **IhorK**  1 year, 9 months ago

Selected Answer: C

"Cluster cache management is active on an Aurora PostgreSQL DB cluster when the cluster has an Aurora Reader instance configured as follows: The Aurora Reader instance uses same DB instance class type and size as the cluster's Writer instance. The Aurora Reader instance is configured as Tier-0 for the cluster. If the cluster has more than one Reader, this is its only Tier-0 Reader. Setting more than one Reader to Tier-0 disables CCM."

https://docs.amazonaws.cn/en_us/AmazonRDS/latest/AuroraUserGuide/aurora_ccm_status.html
upvoted 1 times

  **jitesh_k** 1 year, 6 months ago

Why do we need to set promotion tier for writer instance? It is already being used and failover will occur when it fails. When writer instance fails, reader replica will be promoted to writer instance. The writer instance that failed is not going to be removed from cluster - correct?
upvoted 1 times

  **sayed** 2 years, 7 months ago

Selected Answer: C

C

for CCM to work

- 1- you need to have one RR with the same instance class type and size as the writer instance
- 2- you need to set promotion priority to 0 for both writer and RR you need

the details here <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.cluster-cache-mgmt.html>
upvoted 4 times

  **sachin** 2 years, 11 months ago

C is correct.

<https://aws.amazon.com/blogs/database/introduction-to-aurora-postgresql-cluster-cache-management/>

upvoted 1 times

🗨️ 👤 **novice_expert** 3 years ago

Selected Answer: C

One Aurora Replica of same instance class as the primary DB

-> Aurora PostgreSQL DB cluster cache management

-> failover priority to tier-0 for the primary DB instance and one replica with the same instance class

-> failover priority to tier-1 for the other replicas

upvoted 2 times

🗨️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: C

C is correct. Cluster Cache Management needs these three:

- Set One Replica to have same priority as Primary

- Set its instance class same as Primary

upvoted 4 times

🗨️ 👤 **tugboat** 3 years, 3 months ago

C is economical, but priorities don't need to be set as largest replica will always be first target option for failovers

upvoted 2 times

🗨️ 👤 **VPup** 3 years, 3 months ago

Selected Answer: C

CCM can be set to only one replica. That's why we put it to tier-0 to guarantee the fail over to the CCM enabled replica.

upvoted 4 times

🗨️ 👤 **thelad** 3 years, 3 months ago

Answer is D. The replicas must be the same size as the primary when cluster cache management is used.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.cluster-cache-mgmt.html>

upvoted 4 times

🗨️ 👤 **novice_expert** 3 years ago

No all replicas but one where you want to fail over to, so ans is C

upvoted 3 times

🗨️ 👤 **shuraosipov** 3 years, 5 months ago

Selected Answer: C

Answer is C.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.cluster-cache-mgmt.html>

upvoted 2 times

🗨️ 👤 **thelad** 3 years, 3 months ago

From that link - "Cluster cache management requires that the designated reader instance have the same instance class type and size (db.r5.2xlarge or db.r5.xlarge, for example) as the writer"

Therefore both replicas are required to be the same size as the primary. I'm going for Answer D

upvoted 1 times

🗨️ 👤 **novice_expert** 3 years ago

not both, just one is good

upvoted 1 times

🗨️ 👤 **GMartinelli** 3 years, 6 months ago

Selected Answer: C

Option C

upvoted 2 times

🗨️ 👤 **Dip11** 3 years, 6 months ago

Ans is C. Because this doc <https://aws.amazon.com/blogs/database/introduction-to-aurora-postgresql-cluster-cache-management/> says to set priority of primary and one replica to zero. Which is not the case with D.

upvoted 2 times

🗨️ 👤 **shantest1** 3 years, 7 months ago

C Answer

Tier Priority 0 is the clue to eliminate D:

upvoted 2 times

  **anon9002** 3 years, 7 months ago

The priority of 0 has nothing to do with it since more than one replica can have the same priority. The only difference appears to be cost and C will be cheaper.

"You can customize the order in which your Aurora Replicas are promoted to the primary instance after a failure by assigning each replica a priority. Priorities range from 0 for the first priority to 15 for the last priority. If the primary instance fails, Amazon RDS promotes the Aurora Replica with the better priority to the new primary instance. You can modify the priority of an Aurora Replica at any time. Modifying the priority doesn't trigger a failover.

More than one Aurora Replica can share the same priority, resulting in promotion tiers. If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size, then Amazon RDS promotes an arbitrary replica in the same promotion tier. "

upvoted 2 times

  **im_not_robot** 2 years, 3 months ago

CCM requires one replica has tier 0.

upvoted 1 times

  **LMax** 3 years, 7 months ago

C and D look equally good to address the problem, but D would cost more as you upgrade all read replicas, not just 1. So for cost saving reasons would go with Answer C.

upvoted 1 times

  **Windy** 3 years, 7 months ago



C for me.

upvoted 1 times

  **myutran** 3 years, 7 months ago

Ans: C

upvoted 1 times

  **JobinAkaJoe** 3 years, 7 months ago

I am torn between C & D.

Performance issue after failover lasts only for a short duration which means it has more to do with cache management than instance sizing. CCM definitely is the solution.

Between C & D, I will go with C considering the fact that originally read-replicas were of lower configuration for cost-saving, so its ideal to have one read-replica matching primary instance size and others with lower configuration if cost is a concern.

upvoted 1 times

A company has a database monitoring solution that uses Amazon CloudWatch for its Amazon RDS for SQL Server environment. The cause of a recent spike in CPU utilization was not determined using the standard metrics that were collected. The CPU spike caused the application to perform poorly, impacting users. A Database Specialist needs to determine what caused the CPU spike. Which combination of steps should be taken to provide more visibility into the processes and queries running during an increase in CPU load? (Choose two.)

- A. Enable Amazon CloudWatch Events and view the incoming T-SQL statements causing the CPU to spike.
- B. Enable Enhanced Monitoring metrics to view CPU utilization at the RDS SQL Server DB instance level.
- C. Implement a caching layer to help with repeated queries on the RDS SQL Server DB instance.
- D. Use Amazon QuickSight to view the SQL statement being run.
- E. Enable Amazon RDS Performance Insights to view the database load and filter the load by waits, SQL statements, hosts, or users.

Suggested Answer: BE

Community vote distribution

BE (100%)

🗳️ 👤 **BillyMadison** Highly Voted 3 years, 7 months ago

B&E as well

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-instance-high-cpu/>

"Several factors can cause an increase in CPU utilization. For example, user-initiated heavy workloads, analytic queries, prolonged deadlocks and lock waits, multiple concurrent transactions, long-running transactions, or other processes that utilize CPU resources.

First, you can identify the source of the CPU usage by:

Using Enhanced Monitoring

Using Performance Insights"

upvoted 11 times

🗳️ 👤 **LMax** Highly Voted 3 years, 7 months ago

B & E for me too

upvoted 7 times

🗳️ 👤 **novice_expert** Most Recent 3 years ago

Selected Answer: BE

you can identify the source of the CPU usage by:

Using Enhanced Monitoring

Using Performance Insights"

upvoted 3 times

🗳️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: BE

Gets good visibility

upvoted 3 times

🗳️ 👤 **myutran** 3 years, 7 months ago

Ans: BE

upvoted 1 times

🗳️ 👤 **Ebi** 3 years, 7 months ago

Answer is B,E

upvoted 2 times

🗳️ 👤 **BillyC** 3 years, 7 months ago

Yes B and E here

upvoted 2 times

🗳️ 👤 **Mickysingh** 3 years, 8 months ago

B and E

upvoted 4 times

A company is using Amazon with Aurora Replicas for read-only workload scaling. A Database Specialist needs to split up two read-only applications so each application always connects to a dedicated replica. The Database Specialist wants to implement load balancing and high availability for the read-only applications.

Which solution meets these requirements?

- A. Use a specific instance endpoint for each replica and add the instance endpoint to each read-only application connection string.
- B. Use reader endpoints for both the read-only workload applications.
- C. Use a reader endpoint for one read-only application and use an instance endpoint for the other read-only application.
- D. Use custom endpoints for the two read-only applications.

Suggested Answer: B

Reference:

<https://rimzy.net/category/amazon-rds/page/4/>

Community vote distribution

D (100%)

🗳️ 👤 **SachinGoel** 2 years, 4 months ago

Selected Answer: D

D is answer

upvoted 1 times

🗳️ 👤 **sachin** 2 years, 10 months ago

Custom endpoints are usually used when custom load balancing is needed, i.e. when specific read application should connect to custom read replica directed by custom rules in load balancing. When all RR are of same size, so need for custom endpoint reader endpoint will do.

upvoted 2 times

🗳️ 👤 **sachin** 2 years, 10 months ago

So B is correct

upvoted 2 times

🗳️ 👤 **ryuhei** 2 years, 11 months ago

Selected Answer: D

Answer: D

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: D

A custom endpoint for an Aurora cluster represents a set of DB instances (here replicas) that you choose. When you connect to the endpoint, Aurora performs load balancing and chooses one of the instances in the group to handle the connection.

we can create like two custom endpoints:

1. With two specific read replicas for application A, another three read replicas for application B.

HA and meet the requirements that A and B will connect to different replicas

upvoted 2 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: D

Definitely (D)

upvoted 2 times

🗳️ 👤 **guru_ji** 3 years, 6 months ago

Answer is: D

upvoted 1 times

🗳️ 👤 **ChauPhan** 3 years, 6 months ago

A is incorrect because it is not HA: "Use a specific instance endpoint for each replica"

=> if the specific instance is down, connection is lost

B is incorrect because the demand requires 02 applications connect to different RO replicas. In case B, the traffic will be distributed among read replicas.



C is same with A, using one endpoint or instance does not meet the HA requirement.

D. Use custom endpoints for the two read-only applications.

D is correct because we can create like two customer endpoints: 1. With two specific read replicas for application A, another three read replicas for application B.

HA and meet the requirements that A and B will connect to different replicas

upvoted 2 times

  **Dip11** 3 years, 6 months ago

Ans is D. In order for application to connect to specific instance, and also provide HA by flipping DNS in case of one read replica goes down.

upvoted 3 times

  **LMax** 3 years, 6 months ago

No doubts that D is the answer.

upvoted 4 times

  **myutran** 3 years, 6 months ago

Ans: D

upvoted 2 times

  **Exia** 3 years, 7 months ago

A. For example, your client application might require more fine-grained load balancing based on workload type. In this case, you can configure multiple clients to connect to different Aurora Replicas in a DB cluster to distribute read workloads.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Overview.Endpoints.html#Aurora.Endpoints.Viewing>

upvoted 1 times

  **Exia** 3 years, 7 months ago

The reader endpoint load-balances connections to available Aurora Replicas in an Aurora DB cluster. It doesn't load-balance individual queries. If you want to load-balance each query to distribute the read workload for a DB cluster, open a new connection to the reader endpoint for each query.

upvoted 1 times

  **JobinAkaJoe** 3 years, 7 months ago



D - Custom endpoints allows distributing workload in a customized way by grouping read-replicas of various size, configuration and location

upvoted 1 times

  **Billhardy** 3 years, 7 months ago

D, custom endpoints

upvoted 2 times

  **Ashoks** 3 years, 7 months ago

D. Custom endpoint for custom requirements

upvoted 2 times

  **halol** 3 years, 7 months ago

the answer is D , dedicated endpoint for each application by custom endpoint


or each application has to use reader endpoint not cluster reader and this option does not exists in the answers.

upvoted 1 times

  **awscamus** 3 years, 8 months ago

D is the answer. Dedicated replica is the key

upvoted 2 times

  **Ebi** 3 years, 8 months ago

Answer is D.

B is not correct, reader endpoint of the cluster performs load balancing across all reader instances, questions say reader instances must be split up between to applications.

upvoted 1 times

An online gaming company is planning to launch a new game with Amazon DynamoDB as its data store. The database should be designated to support the following use cases:

- ⇒ Update scores in real time whenever a player is playing the game.
- ⇒ Retrieve a player's score details for a specific game session.

A Database Specialist decides to implement a DynamoDB table. Each player has a unique `user_id` and each game has a unique `game_id`.

Which choice of keys is recommended for the DynamoDB table?

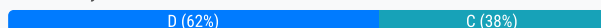
- A. Create a global secondary index with `game_id` as the partition key
- B. Create a global secondary index with `user_id` as the partition key
- C. Create a composite primary key with `game_id` as the partition key and `user_id` as the sort key
- D. Create a composite primary key with `user_id` as the partition key and `game_id` as the sort key

Suggested Answer: B

Reference:

<https://aws.amazon.com/blogs/database/amazon-dynamodb-gaming-use-cases-and-design-patterns/>

Community vote distribution



🗳️ **BillyMadison** Highly Voted 3 years, 8 months ago

I'm going with D based on the following links.

<https://aws.amazon.com/blogs/database/amazon-dynamodb-gaming-use-cases-and-design-patterns/>

"EA uses the user ID as the partition key and primary key (a 1:1 modeling pattern)."

<https://aws.amazon.com/blogs/database/choosing-the-right-dynamodb-partition-key/>

"Partition key and sort key: Referred to as a composite primary key, this type of key is composed of two attributes. The first attribute is the partition key, and the second attribute is the sort key."

upvoted 19 times

🗳️ **JobinAkaJoe** Highly Voted 3 years, 7 months ago

D is the best choice of keys for below requirements.

- ⇒ Update scores in real time whenever a player is playing the game. -- `user_id` being partition key
- ⇒ Retrieve a player's score details for a specific game session. -- `game_id` sort key

upvoted 10 times

🗳️ **MultiAZ** Most Recent 1 year, 4 months ago

Selected Answer: D

I go for D

- `user_ID` for global score

- `user_ID` (PK) + `game_id` (SK) for a session

upvoted 1 times

🗳️ **silvaa360** 1 year, 5 months ago

Selected Answer: C

- `game_id` must reflect the game session and not the actual game, because if not even `user_id` + `game_id` would not be unique. Saying that, `game_id` will have higher cardinality.

- about the query... if it is to get the details of one specific player, both c and d will work, if we want to get the details of all players for a game, only c works.

Having this said, for me C) is the only answer that will work for all cases, having a higher cardinality value as partition key.

But let me say that these types of questions must be better elaborated, because at some times it depends on the interpretation

upvoted 3 times

🗳️ **dnelub** 1 year, 11 months ago

this is my humble contribution to the solution:

If we want to get the high score for a player for a certain game, `user_id` should be our partition key and `game_id` should be our sort key, which is D, but to find out what the highest score ever was for a certain game, we would need `game_id` to be a partition key and 'score' to be the sort key, which is not in the options.

I mean that cardinality can be decided according to the objective of the query but it is not given explicitly in the question.

upvoted 1 times

🗨️ 👤 **jthuma** 2 years, 5 months ago

D. `Game_id` would have lower cardinality than `user_id`. Say the gaming company only had 4 games with 2m users. You would want your partition to be by higher cardinality and your secondary with lower.

upvoted 2 times

🗨️ 👤 **Paulv82003** 1 year, 12 months ago

"each game has a unique `game_id`" what are you talking about "`Game_id` would have lower cardinality than `user_id`"? both unique.

upvoted 2 times

🗨️ 👤 **fserrano** 2 years, 6 months ago

I'm going with C, for queries and updates `game_id` and `user_id` are always used, in addition, `game_id` has more high cardinality and avoids "hot" partitions if used as the partition key.

upvoted 3 times

🗨️ 👤 **JeanGat** 2 years, 7 months ago

Selected Answer: D

D. There are going to be more `user_ids` than `game_ids`, hence `user_id` will be the higher cardinality value = the better partition key. For those of you saying "C", are you telling me Electronic Arts has more unique games than they have users!!??!

<https://aws.amazon.com/blogs/database/choosing-the-right-dynamodb-partition-key/>

upvoted 5 times

🗨️ 👤 **TiredDad** 1 year, 11 months ago

each game has a unique `game_id`.. meaning if you play same game (e.g. car racing) two times, each time you play the game, it will have a unique `game_id`.. that is what I understand from the question.. because you can't have just one game id for car racing then how will you store data for same user id playing same game multiple times?

upvoted 1 times

🗨️ 👤 **praffuln** 3 years ago

Selected Answer: C

If there will be `player_id` it will cause of hot partitioning. So D is wrong.

C is correct as `game_id` will have more high cardinality

upvoted 2 times

🗨️ 👤 **khchan123** 3 years ago

C is the answer.

⇒ Scores should be updated in real time anytime a player is engaged in the game. -> `game_id` is the partition key

⇒ Retrieve the information of a player's score for a certain gaming session. -> `game_id` and `user_id` is the primary key (i.e. `user_id` is the sort key)

upvoted 3 times

🗨️ 👤 **novice_expert** 3 years ago

Selected Answer: D

use cases:

⇒ Scores should be updated in real time anytime a player is engaged in the game. (query based on `user_id` only => `user_id` partition key)

⇒ Retrieve the information of a player's score for a certain gaming session. (query based on `user_id` + `game_id` => `user_id` partition key, and `game_id` sort key)

D. Create a composite primary key with `user_id` as the partition key and `game_id` as the sort key

upvoted 2 times

🗨️ 👤 **TiredDad** 1 year, 11 months ago

Scores should be updated in real time anytime a player is engaged in the game. in this case, query will be based on `user_id` and `game_id`..

otherwise you will end up fetching all records for that `user_id`

upvoted 1 times

🗨️ 👤 **randss** 3 years, 1 month ago

Answer is C, partition key must have high cardinality ie. game_id as there may be multiple instances of games played by each user.
upvoted 2 times

🗨️ 👤 **ugo009** 3 years, 2 months ago

I chose D. GSI does not support strong consistency.
upvoted 1 times

🗨️ 👤 **ChauPhan** 3 years, 7 months ago

For me, both C and D can be used but C is better.

⇒ Update scores in real time whenever a player is playing the game.

⇒ Query by the game_id first, then update relevant user_id with his relevant score.

⇒ Retrieve a player's score details for a specific game session => game_id can be used to query then specify the user_id

D can do the same as above. But for the query such as: Top score player for each game, C is better.

D is for the query such as: This player plays how many games and the score of each.

upvoted 1 times

🗨️ 👤 **uupadhyay** 3 years, 7 months ago

Why not C ? Any idea ?

upvoted 1 times

🗨️ 👤 **jove** 3 years, 5 months ago

Partition key should have high cardinality. Therefore user id is better than game id as the partition key.. Answer is D

upvoted 5 times

🗨️ 👤 **Dip11** 3 years, 7 months ago

I would go with A.

Reason, Need user_id as partition key for main table and global secondary index with game_id as partition key since we want to update via user_id and also retrieve via game_id. Both these options do not exist in any of the options so A seems to be nearest correct answer.

upvoted 2 times

🗨️ 👤 **nideesh** 3 years, 5 months ago

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

GameScores is identified by a partition key (UserId) and a sort key (GameTitle). A query that specified the key attributes (UserId and GameTitle) would be very efficient

Composite primary key is the answer. For building a leader board we need global secondary index

upvoted 1 times

🗨️ 👤 **LMax** 3 years, 7 months ago

Answer D

upvoted 3 times

A Database Specialist migrated an existing production MySQL database from on-premises to an Amazon RDS for MySQL DB instance. However, after the migration, the database needed to be encrypted at rest using AWS KMS. Due to the size of the database, reloading the data into an encrypted database would be too time-consuming, so it is not an option.

How should the Database Specialist satisfy this new requirement?

- A. Create a snapshot of the unencrypted RDS DB instance. Create an encrypted copy of the unencrypted snapshot. Restore the encrypted snapshot copy.
- B. Modify the RDS DB instance. Enable the AWS KMS encryption option that leverages the AWS CLI.
- C. Restore an unencrypted snapshot into a MySQL RDS DB instance that is encrypted.
- D. Create an encrypted read replica of the RDS DB instance. Promote it the master.

Suggested Answer: A

Community vote distribution

A (100%)

 **BillyMadison** Highly Voted 3 years, 8 months ago

Agree with A. This blog post also says the same steps as Micky's link.

"However, because you can encrypt a copy of an unencrypted DB snapshot, you can effectively add encryption to an unencrypted DB instance. That is, you can create a snapshot of your DB instance, and then create an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot, and thus you have an encrypted copy of your original DB instance. For more information, see Copying a Snapshot."

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html>

upvoted 16 times

 **IhorK** Most Recent 1 year, 9 months ago

Selected Answer: A

Encrypt your existing AWS RDS:

- Take snapshot
- Copy Snapshot with Enable encryption
- Change unencrypted instance name

Encrypt your existing AWS RDS:

- Take snapshot
- Copy Snapshot with Enable encryption
- Change unencrypted instance name
- Restore Snapshot (DB Instance Identifier field must contain the name of your previous database before renaming it)

<https://blog.theodo.com/2019/11/encrypt-existing-aws-rds-database/>

upvoted 1 times

 **novice_expert** 3 years ago

Selected Answer: A

snapshot -> encrypted snapshot -> Restore the encrypted snapshot copy.

upvoted 2 times

 **Mintwater** 2 years, 1 month ago

Limitations of Amazon RDS encrypted DB instances

The following limitations exist for Amazon RDS encrypted DB instances:

You can only encrypt an Amazon RDS DB instance when you create it, not after the DB instance is created.

However, because you can encrypt a copy of an unencrypted snapshot, you can effectively add encryption to an unencrypted DB instance. That is, you can create a snapshot of your DB instance, and then create an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot, and thus you have an encrypted copy of your original DB instance. For more information, see Copying a DB snapshot.

You can't turn off encryption on an encrypted DB instance.

You can't create an encrypted snapshot of an unencrypted DB instance.

upvoted 4 times

🗳️ 👤 **kret** 3 years, 1 month ago

Selected Answer: A

A is an obvious choice here

upvoted 2 times

🗳️ 👤 **uupadhyay** 3 years, 6 months ago

Here source database is in on premises ? How we can access snapshot of on-premises database from AWS console ? How A is possible ?

upvoted 1 times

🗳️ 👤 **Scunningham99** 3 years, 6 months ago

its talking about the database has already been migrated to RDS, but they want to know how to encrypt it after the migration - hence a is correct

upvoted 2 times

🗳️ 👤 **LMax** 3 years, 6 months ago

Answer A

upvoted 2 times

🗳️ 👤 **myutran** 3 years, 6 months ago

Ans: A

upvoted 1 times

🗳️ 👤 **Roontha** 3 years, 7 months ago

Answer: A

upvoted 1 times

🗳️ 👤 **ppravesh** 3 years, 7 months ago

I guess It should be D: As per this document it is possible to do replication from unencrypted to encrypted read replica.

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-encrypt-instance-mysql-mariadb/>

upvoted 1 times

🗳️ 👤 **GeeBeeEl** 3 years, 7 months ago

This is for MariaDB!!!

upvoted 1 times

🗳️ 👤 **Exia** 3 years, 7 months ago

You cannot create an encrypted Read Replica from an unencrypted DB instance.

upvoted 1 times

🗳️ 👤 **rootkim** 3 years, 8 months ago

If you become D, do you really need to do A?

How do you fix data inconsistencies?

D is possible, so I think it's best.

upvoted 1 times

🗳️ 👤 **JobinAkaJoe** 3 years, 7 months ago

To have an encrypted read-replica, primary instance must be encrypted

upvoted 2 times

🗳️ 👤 **BillyC** 3 years, 8 months ago

A is correct

upvoted 1 times

🗳️ 👤 **Mickysingh** 3 years, 8 months ago

A is correct answer

<https://blog.theodo.com/2019/11/encrypt-existing-aws-rds-database/>

upvoted 4 times

A Database Specialist is planning to create a read replica of an existing Amazon RDS for MySQL Multi-AZ DB instance. When using the AWS Management

Console to conduct this task, the Database Specialist discovers that the source RDS DB instance does not appear in the read replica source selection box, so the read replica cannot be created.

What is the most likely reason for this?

- A. The source DB instance has to be converted to Single-AZ first to create a read replica from it.
- B. Enhanced Monitoring is not enabled on the source DB instance.
- C. The minor MySQL version in the source DB instance does not support read replicas.
- D. Automated backups are not enabled on the source DB instance.

Suggested Answer: D

Reference:

<https://aws.amazon.com/rds/features/read-replicas/>

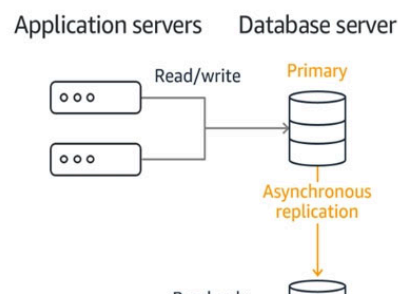
Amazon RDS Read Replicas

Amazon RDS Read Replicas provide enhanced performance and durability for RDS database (DB) instances. They make it easy to elastically scale out beyond the capacity constraints of a single DB instance for read-heavy database workloads. You can create one or more replicas of a given source DB Instance and serve high-volume application read traffic from multiple copies of your data, thereby increasing aggregate read throughput. Read replicas can also be promoted when needed to become standalone DB instances. Read replicas are available in Amazon RDS for [MySQL](#), [MariaDB](#), [PostgreSQL](#), [Oracle](#), and [SQL Server](#) as well as [Amazon Aurora](#).

For the MySQL, MariaDB, PostgreSQL, Oracle, and SQL Server database engines, Amazon RDS creates a second DB instance using a snapshot of the source DB instance. It then uses the engines' native asynchronous replication to update the read replica whenever there is a change to the source DB instance. The read replica operates as a DB instance that allows only read-only connections; applications can connect to a read replica just as they would to any DB instance. Amazon RDS

Community vote distribution

D (100%)



RotterDam Highly Voted 3 years, 2 months ago

Selected Answer: D

(D) is correct

In order to create a read replica

- Automated Backups MUST be enabled

Additionally for SQL Server

- Source MUST be Multi-AZ deployment with Always On Availability Groups (AG)

upvoted 6 times

novice_expert Highly Voted 3 years ago

Selected Answer: D

D. Automated backups are not enabled on the source DB instance.

In order to create a read replica

- Automated Backups MUST be enabled

Additionally for SQL Server

- Source MUST be Multi-AZ deployment with Always On Availability Groups (AG)

upvoted 5 times

Pranava_GCP Most Recent 1 year, 8 months ago



Selected Answer: D

D. Automated backups are not enabled on the source DB instance.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_MultiAZDBCluster_ReadRepl.html



"You must turn on automatic backups on the source DB instance by setting the backup retention period to a value other than 0."

upvoted 1 times

  **awsmonster** 3 years, 4 months ago

Ans: D

upvoted 2 times

  **avland** 3 years, 6 months ago

D.

>Your source DB instance must have backup retention enabled.

https://docs.aws.amazon.com/AmazonRDS/latest/APIReference/API_CreateDBInstanceReadReplica.html

upvoted 2 times

A Database Specialist has migrated an on-premises Oracle database to Amazon Aurora PostgreSQL. The schema and the data have been migrated successfully.

The on-premises database server was also being used to run database maintenance cron jobs written in Python to perform tasks including data purging and generating data exports. The logs for these jobs show that, most of the time, the jobs completed within 5 minutes, but a few jobs took up to 10 minutes to complete. These maintenance jobs need to be set up for Aurora PostgreSQL.

How can the Database Specialist schedule these jobs so the setup requires minimal maintenance and provides high availability?

- A. Create cron jobs on an Amazon EC2 instance to run the maintenance jobs following the required schedule.
- B. Connect to the Aurora host and create cron jobs to run the maintenance jobs following the required schedule.
- C. Create AWS Lambda functions to run the maintenance jobs and schedule them with Amazon CloudWatch Events.
- D. Create the maintenance job using the Amazon CloudWatch job scheduling plugin.


Suggested Answer: D

Reference:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/mw-cli-task-options.html>

Community vote distribution

C (100%)

  **Mickysingh**  3 years, 8 months ago

Answer should be C but below link confuses

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/Create-CloudWatch-Events-Scheduled-Rule.html>

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/schedule-jobs-for-amazon-rds-and-aurora-postgresql-using-lambda-and-secrets-manager.html>

a job for data extraction or a job for data purging can easily be scheduled using cron. For these jobs, database credentials are typically either hard-coded or stored in a properties file. However, when you migrate to Amazon Relational Database Service (Amazon RDS) or Amazon Aurora PostgreSQL, you lose the ability to log in to the host instance to schedule cron jobs.

This pattern describes how to use AWS Lambda and AWS Secrets Manager to schedule jobs for Amazon RDS and Aurora PostgreSQL databases after migration.

it confirms that answer is C

upvoted 13 times

  **BillyMadison** 3 years, 8 months ago

I agree with C via your second link:

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/schedule-jobs-for-amazon-rds-and-aurora-postgresql-using-lambda-and-secrets-manager.html>

upvoted 3 times

  **waterh30** 3 years, 8 months ago

<https://medium.com/better-programming/cron-job-patterns-in-aws-126fbf54a276>



upvoted 2 times

  **Pranava_GCP**  1 year, 8 months ago

Selected Answer: C

C. Create AWS Lambda functions to run the maintenance jobs and schedule them with Amazon CloudWatch Events.

upvoted 1 times

  **IhorK** 1 year, 9 months ago

Selected Answer: C

- You can also use CloudWatch Events to schedule automated actions that self-initiate at certain times using cron or rate expressions.
- Lambda supports Python.
- CloudWatch Events can call Lambda functions.
- Less than 15 min job.

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/schedule-jobs-for-amazon-rds-for-postgresql-and-aurora-postgresql-by-using-lambda-and-secrets-manager.html>

upvoted 2 times

🗨️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: C

C. Whenever its asked to create Cron Jobs Lambda vs EC2 : Its ALWAYS Lambda!!

upvoted 4 times

🗨️ 👤 **im_not_robot** 2 years, 3 months ago

It depends on how long the job run. Because lambda has 15 minutes limitation.

If the job needs more than 15 minutes to run, then can consider AWS Batch, ECS, EC2

upvoted 2 times

🗨️ 👤 **GMartinelli** 3 years, 5 months ago

Selected Answer: C

Option C

upvoted 1 times

🗨️ 👤 **guru_ji** 3 years, 6 months ago

Correct Answer ==>> C

upvoted 2 times

🗨️ 👤 **ChauPhan** 3 years, 6 months ago

C is correct.

A B is not HA.

D, there is no such maintenance JOB for CW Events. CW Events has to trigger the other functions (Lambda/SSM/Step Function/ Batch) to do the task.

upvoted 1 times

🗨️ 👤 **guru_ji** 3 years, 6 months ago

Correct Answer: C

upvoted 1 times

🗨️ 👤 **Dip11** 3 years, 7 months ago

Answer should be C. I don't think there is anything like CloudWatch job scheduling plugin. Internet search doesn't show anything like this.

upvoted 2 times

🗨️ 👤 **LMax** 3 years, 7 months ago

Agree with Answer C

upvoted 2 times

🗨️ 👤 **LMax** 3 years, 7 months ago

Agree with Answer C

upvoted 2 times

🗨️ 👤 **Windy** 3 years, 7 months ago

C for me

upvoted 1 times

🗨️ 👤 **myutran** 3 years, 7 months ago

Ans: C

upvoted 2 times

🗨️ 👤 **JobinAkaJoe** 3 years, 7 months ago

C is the best answer.

Eventbridge would have been a better choice to schedule lambda.

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/RunLambdaSchedule.html>

upvoted 2 times

🗨️ 👤 **Ashoks** 3 years, 8 months ago

Ans is C. CloudWatch scheduling and Lambda execution and this option should be fine as long as job completes within 15 minutes.

upvoted 4 times

🗨️ 👤 **BillyC** 3 years, 8 months ago

Answer should be C

upvoted 4 times

  **[Removed]** 3 years, 8 months ago

The document referenced has no mention of job scheduling plugin for cloudwatch.... it is for systems manager.

I will have to go with C.

upvoted 2 times

A company has an Amazon RDS Multi-AZ DB instances that is 200 GB in size with an RPO of 6 hours. To meet the company's disaster recovery policies, the database backup needs to be copied into another Region. The company requires the solution to be cost-effective and operationally efficient.

What should a Database Specialist do to copy the database backup into a different Region?

- A. Use Amazon RDS automated snapshots and use AWS Lambda to copy the snapshot into another Region
- B. Use Amazon RDS automated snapshots every 6 hours and use Amazon S3 cross-Region replication to copy the snapshot into another Region
- C. Create an AWS Lambda function to take an Amazon RDS snapshot every 6 hours and use a second Lambda function to copy the snapshot into another Region
- D. Create a cross-Region read replica for Amazon RDS in another Region and take an automated snapshot of the read replica

Suggested Answer: D

Reference:

<https://aws.amazon.com/blogs/database/implementing-a-disaster-recovery-strategy-with-amazon-rds/>

Community vote distribution

C (100%)

  **Ebi** Highly Voted 3 years, 7 months ago

Answer is C

A. Use Amazon RDS automated snapshots and use AWS Lambda to copy the snapshot into another Region

Automated snapshots are taken once per day only, RPO is 6 hours, so not an option

B. Use Amazon RDS automated snapshots every 6 hours and use Amazon S3 cross-Region replication to copy the snapshot into another Region

You can not take automated snapshots every 6 hours

C. Create an AWS Lambda function to take an Amazon RDS snapshot every 6 hours and use a second Lambda function to copy the snapshot into another Region

Only possible option

D. Create a cross-Region read replica for Amazon RDS in another Region and take an automated snapshot of the read replica

Not cost-effective, replica is the most expensive DR option.

upvoted 27 times

  **szmulder** 3 years, 7 months ago

It's hard to choose. A is correct is RDS can restore to point in time, so we don't need to do a snapshot every 6 hours.



upvoted 1 times

  **zanhsieh** 3 years, 7 months ago

Agreed with C. Any options mention / use 'automated snapshot' should be dropped, so no ABD. Verified via AWS console in RDS. Change time interval for RDS automatic backup period to 0 means disable automatic backup. See:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithAutomatedBackups.html

upvoted 3 times

  **MultiAZ** Most Recent 1 year, 4 months ago

Selected Answer: C

Answer is C.

D is not the most cost-effective

A and B are not OK because automated snapshots are only once a day and you need one every 6 hours

upvoted 1 times

  **SuriSagar** 1 year, 7 months ago

C is the correct answer

upvoted 2 times

🗨️ **aviathor** 1 year, 12 months ago

It would be so much easier to use cross-region backup replication...

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReplicateBackups.html

upvoted 1 times

🗨️ **sachin** 2 years, 10 months ago

Copying 200 Gb of snapshot every 6 hours across region will be costlier approach and ensuring lambda to finish job of copying 200gb snapshot with 15 mins is also not viable.

So i think D is best choice.

upvoted 3 times

🗨️ **awsguys** 3 years ago

c is right .

upvoted 1 times

🗨️ **novice_expert** 3 years ago

Selected Answer: C

x A & B out: only automated snapshot per day allowed we need 4

C. Create an AWS Lambda function to take an Amazon RDS snapshot every 6 hours and use a second Lambda function to copy the snapshot into another Region (cost effective)

D. Create a cross-Region read replica for Amazon RDS in another Region and take an automated snapshot of the read replica (costly)

upvoted 1 times

🗨️ **RotterDam** 3 years, 2 months ago

Selected Answer: C

Definitely (C) - also a question in the official sample exam discussed by Stephen Maarek and Riyaz in their Udemy course

upvoted 1 times

🗨️ **Shunpin** 3 years, 5 months ago

Selected Answer: C

System snapshot can't fulfill 6 hours requirement. You need to control it by script

<https://aws.amazon.com/blogs/database/%C2%AD%C2%AD%C2%ADautomating-cross-region-cross-account-snapshot-copies-with-the-snapshot-tool-for-amazon-aurora/>

upvoted 1 times

🗨️ **guru_ji** 3 years, 6 months ago

Anyone planning for exam?

We can share study material with each other, it would be beneficial for both. You can email me on "awsdbguru at gmail"

upvoted 1 times

🗨️ **guru_ji** 3 years, 6 months ago

Correct Answer is ==>> C. any idea how much Q we will get in real exam from Q available here? anyone is preparing for this exam and want to do group study with us, comment with mail id.

upvoted 1 times

🗨️ **ChauPhan** 3 years, 6 months ago

<https://aws.amazon.com/blogs/database/implementing-a-disaster-recovery-strategy-with-amazon-rds/>

Automated Backup is daily, so 6 hour RPO is not possible.

Only C is correct.

upvoted 1 times

🗨️ **Dip11** 3 years, 6 months ago

Answer is C. Reason for eliminating D is that RDS SQL SERVER does not support cross-region read replica and the question does not state which database engine used in RDS.

upvoted 1 times

🗨️ **LMax** 3 years, 6 months ago

Agree with C

upvoted 2 times

🗨️ **Windy** 3 years, 7 months ago

I think it is C.

upvoted 1 times

🗨️ 👤 **myutran** 3 years, 7 months ago

Ans: D with key "operationally efficient".

With option C, it is necessary to copy 200GB of data from one region to another every 6 hours.

upvoted 2 times

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

With C, it would be required for a Lambda function to copy 200 GB in 15 mins (max execution time) - every 6 hours. It might be feasible given a decent throughput (cross-region!), but that's not guaranteed and there's no mention of the available throughput.

upvoted 1 times

🗨️ 👤 **JobinAkaJoe** 3 years, 7 months ago

A - wrong - I guess you cannot copy automated snapshots to another region. You must create a copy of automated snapshot in the same region first.

B - wrong - automated snapshots are taken once in a day which alone wouldnt meet proposed RPO of 6hrs

C - best answer for the requirement in hand.

D - best option for cross region DR. However having a read-replica in another region is costlier than option C

upvoted 1 times

An Amazon RDS EBS-optimized instance with Provisioned IOPS (PIOPS) storage is using less than half of its allocated IOPS over the course of several hours under constant load. The RDS instance exhibits multi-second read and write latency, and uses all of its maximum bandwidth for read throughput, yet the instance uses less than half of its CPU and RAM resources.


What should a Database Specialist do in this situation to increase performance and return latency to sub-second levels?

- A. Increase the size of the DB instance storage
- B. Change the underlying EBS storage type to General Purpose SSD (gp2)
- C. Disable EBS optimization on the DB instance
- D. Change the DB instance to an instance class with a higher maximum bandwidth

Suggested Answer: B

Community vote distribution

D (100%)


 **BillyMadison** Highly Voted 3 years, 8 months ago

I think this is D

https://docs.amazonaws.cn/en_us/AmazonRDS/latest/UserGuide/CHAP_BestPractices.html

"If you are already using Provisioned IOPS storage, provision additional throughput capacity." Does D sound right?

upvoted 13 times

 **MultiAZ** Most Recent 1 year, 4 months ago

Selected Answer: D

Definitely D

A will not help with the bottleneck, B and C will do things worse

upvoted 1 times

 **Pranava_GCP** 1 year, 8 months ago

Selected Answer: D

D. Change the DB instance to an instance class with a higher maximum bandwidth

the performance issue in this case is not related to storage nor cpu nor RAM, it is related to network bandwidth.

upvoted 2 times

 **IhorK** 1 year, 9 months ago

Selected Answer: D

We need to understand the difference between throughput and IOPS.

IOPS – Count of read/write operations per second.

Throughput – Count of read/write bits per second (bps). This measures the amount of time it takes for a disk to read and write data. Throughput is typically the best storage metric when measuring data that needs to be streamed rapidly, such as images and video files.

So it's not about the size (A).

Provisioned IOPS storage to General Purpose SSD (gp2) (B) - rather worse performance judging by data from

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html>

Disable EBS optimization (E) - rather, it should be activated, if there is such a thing.

Instance class with a higher maximum bandwidth (D) - we have "uses all of its maximum bandwidth for read throughput" in the question.

We need to choose another instance class with a higher bandwidth - D.

upvoted 1 times

 **Iollyj** 2 years, 5 months ago

Selected Answer: D

Correct answer I think

upvoted 1 times

 **novice_expert** 3 years ago

Selected Answer: D

Objective is to handle maximum bandwidth for read throughput used

- x A. Increase the size of the DB instance storage (unrelated)
- x B. Change the underlying EBS storage type to General Purpose SSD (gp2) (will slow down)
- x C. Disable EBS optimization on the DB instance (will slow down)
- D. Change the DB instance to an instance class with a higher maximum bandwidth

https://docs.amazonaws.cn/en_us/AmazonRDS/latest/UserGuide/CHAP_BestPractices.html

upvoted 2 times

🗳️ 👤 **kret** 3 years, 1 month ago

Selected Answer: D

- A. Increase the size of the DB instance storage - nonsense; instance is using EBS
- B. Change the underlying EBS storage type to General Purpose SSD (gp2) -> nonsense; GP2 is slower than IO2
- C. Disable EBS optimization on the DB instance -> nonsense; nothing like this exist

D is the answer

upvoted 3 times

🗳️ 👤 **guru_ji** 3 years, 7 months ago

Correct Answer ==>> D

upvoted 1 times

🗳️ 👤 **LMax** 3 years, 7 months ago

My answer is D

upvoted 1 times

🗳️ 👤 **myutran** 3 years, 7 months ago

Ans: D

upvoted 1 times

🗳️ 👤 **RSSRAO** 3 years, 7 months ago

Answer is D

upvoted 1 times

🗳️ 👤 **JobinAkaJoe** 3 years, 7 months ago

D is the right answer

upvoted 1 times

🗳️ 👤 **Ashoks** 3 years, 7 months ago

D. For high throughput.

upvoted 2 times

🗳️ 👤 **AWSCert2020** 3 years, 7 months ago

D here! Throughput is related to Size/Type of Instance

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-optimized.html>

upvoted 1 times

🗳️ 👤 **AWSCert2020** 3 years, 7 months ago

Throughput is related to Size/Type of Instance

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-optimized.html>

upvoted 1 times

🗳️ 👤 **BillyC** 3 years, 8 months ago

Sorry i mean D

upvoted 2 times

🗳️ 👤 **BillyC** 3 years, 8 months ago

Ans B here

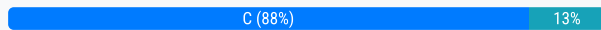
upvoted 1 times

After restoring an Amazon RDS snapshot from 3 days ago, a company's Development team cannot connect to the restored RDS DB instance. What is the likely cause of this problem?

- A. The restored DB instance does not have Enhanced Monitoring enabled
- B. The production DB instance is using a custom parameter group
- C. The restored DB instance is using the default security group
- D. The production DB instance is using a custom option group

Suggested Answer: B

Community vote distribution



🗳️ 👤 **learnaws** Highly Voted 3 years, 7 months ago

C.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RestoreFromSnapshot.html

upvoted 12 times

🗳️ 👤 **BillyMadison** 3 years, 7 months ago

C, agree.

upvoted 3 times

🗳️ 👤 **BillyC** Highly Voted 3 years, 8 months ago

C is correct

upvoted 5 times

🗳️ 👤 **Sathish_dbs** Most Recent 1 year, 7 months ago

is it misleading us to assume that the DB connected for first three days?

upvoted 1 times

🗳️ 👤 **narvaez** 1 year, 7 months ago

Selected Answer: C

C is correct

upvoted 1 times

🗳️ 👤 **saikarthikeya777** 1 year, 11 months ago

Selected Answer: C

how do we know what parameter or option groups source rds is using so assuming security group mismatch because even parameter or option group mismatch also if security group is correct one we can make the connections

upvoted 1 times

🗳️ 👤 **Kodoma** 2 years ago

Selected Answer: C

It is C.

upvoted 1 times

🗳️ 👤 **Chirantan** 2 years, 10 months ago

When you restore a DB instance, the default virtual private cloud (VPC), DB subnet group, and VPC security group are associated with the restored instance, unless you choose different ones.

upvoted 3 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: B

C and B both have incomplete info, would go with B assuming that PROD is the source, and custom parameter group not selected or changed from default for restored one

x A. Enhanced Monitoring enabled (unrelated - for performance check)

B. The production DB instance is using a custom parameter group (would be correct if it was the source)

x C. The restored DB instance is using the default security group (but would not be issue if source also had default)

x D. The production DB instance is using a custom option group (unrelated - it can specify features, called options, that are available for a particular Amazon RDS DB instance)

<https://aws.amazon.com/about-aws/whats-new/2018/10/specify-parameter-groups-when-restoring-amazon-rds-backups/>
upvoted 1 times

🗨️ **novice_expert** 3 years ago

Correction: Ans is C

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RestoreFromSnapshot.html

When you restore a DB instance, the default virtual private cloud (VPC), DB subnet group, and VPC security group are associated with the restored instance, unless you choose different ones.

upvoted 2 times

🗨️ **kret** 3 years, 1 month ago

Selected Answer: C

SG is most typically the cause of an connection issue

upvoted 4 times

🗨️ **soyyodario** 3 years, 3 months ago

C, but B.....

Parameter group considerations

We recommend that you retain the DB parameter group for any DB snapshots you create, so that you can associate your restored DB instance with the correct parameter group.

The default DB parameter group is associated with the restored instance, unless you choose a different one. No custom parameter settings are available in the default parameter group.

You can specify the parameter group when you restore the DB instance.

upvoted 1 times

🗨️ **ChauPhan** 3 years, 6 months ago

Omit B, D regarding PROD DB

A is not relevant.

C is correct.

upvoted 3 times

🗨️ **gelsm** 3 years, 7 months ago

C. The restored DB instance is using the default security group

This is likely the cause of the problem since security groups control the connectivity to the DB instance.

upvoted 2 times

🗨️ **gelsm** 3 years, 6 months ago

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-cannot-connect/>

upvoted 1 times

🗨️ **LMax** 3 years, 7 months ago

Answer C

upvoted 4 times

🗨️ **myutran** 3 years, 7 months ago

Ans: C

upvoted 2 times

🗨️ **kubilay** 3 years, 7 months ago



C is correct

upvoted 2 times

🗨️ **JobinAkaJoe** 3 years, 7 months ago

C is the correct answer

upvoted 2 times

  **Ashoks** 3 years, 7 months ago

yes. C. SG controls the access.

upvoted 3 times

A gaming company has implemented a leaderboard in AWS using a Sorted Set data structure within Amazon ElastiCache for Redis. The ElastiCache cluster has been deployed with cluster mode disabled and has a replication group deployed with two additional replicas. The company is planning for a worldwide gaming event and is anticipating a higher write load than what the current cluster can handle. Which method should a Database Specialist use to scale the ElastiCache cluster ahead of the upcoming event?

- A. Enable cluster mode on the existing ElastiCache cluster and configure separate shards for the Sorted Set across all nodes in the cluster.
- B. Increase the size of the ElastiCache cluster nodes to a larger instance size.
- C. Create an additional ElastiCache cluster and load-balance traffic between the two clusters.
- D. Use the EXPIRE command and set a higher time to live (TTL) after each call to increment a given key.

Suggested Answer: B

Reference:

<https://aws.amazon.com/blogs/database/work-with-cluster-mode-on-amazon-elasticache-for-redis/>

AWS Database Blog

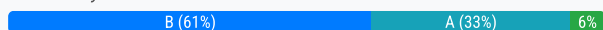
How to work with Cluster Mode on Amazon ElastiCache for Redis

by Josh Kahn | on 26 JUL 2019 | in [Amazon ElastiCache](#), [Database](#) | [Permalink](#) | [Share](#)

One of my favorite aspects of AWS is the large number of building blocks available to tackle a wide range of technical use cases. [Amazon ElastiCache for Redis](#) is one of those building blocks. While most often thought of for fast database caching needs, ElastiCache is quite flexible and fast (microseconds). In the past, I have discussed how to [perform geospatial queries](#) and [build a real-time dashboard](#) with ElastiCache.

In this post, I will describe how you can leverage ElastiCache for Redis with cluster mode enabled to enhance reliability and availability with little change to your existing workload. Cluster Mode comes with the primary benefit of horizontal scaling up and down of your Redis cluster, with almost zero impact on the performance of the cluster, as I will demonstrate later. If you have ever encountered a Redis cluster that is over or under-provisioned or just want to better understand its inner workings, please read on.

Community vote distribution



novice_expert Highly Voted 3 years ago

Selected Answer: B

cannot enable Cluster Mode on an existing cluster, With cluster mode disabled it will allow only vertical scaling.
upvoted 5 times

rrshah83 Most Recent 1 year, 5 months ago

Selected Answer: A

Elasticache now supports enabling cluster mode on existing clusters.

<https://aws.amazon.com/about-aws/whats-new/2023/05/amazon-elasticache-redis-cluster-mode-configuration-existing-clusters/#:~:text=Amazon%20ElastiCache%20for%20Redis%20now%20supports%20enabling%20Cluster%20Mode%20configuration%20on%20existing%20clusters,Posted%20on%3A%20May&text=You%20can%20now%20update%20your,data%2C%20or%20affect%20application%20availability.>

Posted%20on%3A%20May&text=You%20can%20now%20update%20your,data%2C%20or%20affect%20application%20availability.

upvoted 3 times

luckybme 1 year, 7 months ago

Selected Answer: A

Elasticache now supports enabling cluster mode on existing clusters.

<https://aws.amazon.com/about-aws/whats-new/2023/05/amazon-elasticache-redis-cluster-mode-configuration-existing-clusters/#:~:text=Amazon%20ElastiCache%20for%20Redis%20now%20supports%20enabling%20Cluster%20Mode%20configuration%20on%20existing%20clusters,Posted%20on%3A%20May&text=You%20can%20now%20update%20your,data%2C%20or%20affect%20application%20availability.>

Posted%20on%3A%20May&text=You%20can%20now%20update%20your,data%2C%20or%20affect%20application%20availability.

upvoted 3 times

Pranava_GCP 1 year, 8 months ago

Selected Answer: B

B. Increase the size of the ElastiCache cluster nodes to a larger instance size.

<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Scaling.RedisReplGrps.html>

upvoted 1 times

🗳️ 👤 **nehacool29** 1 year, 11 months ago

Selected Answer: C

C is correct

upvoted 1 times

🗳️ 👤 **dnelub** 1 year, 11 months ago

Option B (increasing the size of cluster nodes) can provide some level of scaling, but it may have limitations in terms of the maximum capacity it can handle. Additionally, simply increasing the node size may not fully address the anticipated higher write load during the gaming event.

Option C: Creating an additional ElastiCache cluster and load-balancing traffic between the clusters allows for distributing the write load across multiple clusters, effectively scaling the capacity and handling increased demand. This approach provides horizontal scalability and helps mitigate the potential performance limitations of a single cluster.

Why not C?

upvoted 1 times

🗳️ 👤 **rags1482** 2 years, 7 months ago

B

Redis (cluster mode disabled) supports scaling. You can scale read capacity by adding or deleting replica nodes, or you can scale capacity by scaling up to a larger node type.

upvoted 2 times

🗳️ 👤 **ryuhei** 2 years, 11 months ago

Selected Answer: B

Answer:B

upvoted 1 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: B

B is correct. You cannot enable Cluster Mode on an existing cluster

upvoted 4 times

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

Cluster mode configuration can only be changed from cluster mode disabled to cluster mode enabled. Reverting this configuration is not possible.

<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/modify-cluster-mode.html>

upvoted 1 times

🗳️ 👤 **mayank830** 3 years, 3 months ago

<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Replication.Redis-RedisCluster.html>

Reads v. writes – If the primary load on your cluster is applications reading data, you can scale a Redis (cluster mode disabled) cluster by adding and deleting read replicas. However, there is a maximum of 5 read replicas. If the load on your cluster is write-heavy, you can benefit from the additional write endpoints of a Redis (cluster mode enabled) cluster with multiple shards.

upvoted 1 times

🗳️ 👤 **jeyp12** 3 years, 3 months ago

I think answer is B. With cluster mode disabled it will allow only vertical scaling. So going with B

upvoted 2 times

🗳️ 👤 **2025flakyt** 3 years, 5 months ago

A is the correct answer

Enabling Cluster Mode provides a number of additional benefits in scaling your cluster. In short, it allows you to scale in or out the number of shards (horizontal scaling) versus scaling up or down the node type (vertical scaling). This means that Cluster Mode can scale to very large amounts of storage (potentially 100s of terabytes) across up to 90 shards, whereas a single node can only store as much data in memory as the instance type has capacity for.

<https://aws.amazon.com/blogs/database/work-with-cluster-mode-on-amazon-elasticache-for-redis/>

upvoted 2 times

🗳️ 👤 **VPup** 3 years, 3 months ago

you can not enable the cluster mode on the already running Elasticache cluster. Have to provision a new cluster with the Cluster mode enabled and restore the backed up data from S3.

upvoted 2 times

  **RotterDam** 3 years, 2 months ago

Wrong. You cannot convert an EXISTING Cluster Mode Disabled cluster to Cluster Mode enabled. You need to create a new cluster with cluster mode enabled and warm cache it by loading it with the RDB file backup of the original

upvoted 2 times

An ecommerce company has tasked a Database Specialist with creating a reporting dashboard that visualizes critical business metrics that will be pulled from the core production database running on Amazon Aurora. Data that is read by the dashboard should be available within 100 milliseconds of an update.

The Database Specialist needs to review the current configuration of the Aurora DB cluster and develop a cost-effective solution. The solution needs to accommodate the unpredictable read workload from the reporting dashboard without any impact on the write availability and performance of the DB cluster.

Which solution meets these requirements?

- A. Turn on the serverless option in the DB cluster so it can automatically scale based on demand.
- B. Provision a clone of the existing DB cluster for the new Application team.
- C. Create a separate DB cluster for the new workload, refresh from the source DB cluster, and set up ongoing replication using AWS DMS change data capture (CDC).
- D. Add an automatic scaling policy to the DB cluster to add Aurora Replicas to the cluster based on CPU consumption.

Suggested Answer: A

Community vote distribution

D (92%)

8%

 **novice_expert** Highly Voted 3 years ago

Selected Answer: D

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Replication.html>

replica lag < 100 ms

Option A would take time

upvoted 7 times


 **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: D

D. Add an automatic scaling policy to the DB cluster to add Aurora Replicas to the cluster based on CPU consumption.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Replication.html>


upvoted 3 times

 **manig** 1 year, 11 months ago

agree it's D...there is no 'serverless turn on' option exist

For example, take a scaling policy that uses the predefined average CPU utilization metric. Such a policy can keep CPU utilization at, or close to, a specified percentage of utilization, such as 40 percent.


upvoted 3 times

 **IBANGA007** 2 years, 5 months ago

Selected Answer: C

C. Create a separate DB cluster for the new workload, refresh from the source DB cluster, and set up ongoing replication using AWS DMS change data capture (CDC).

upvoted 1 times

 **lollyj** 2 years, 5 months ago

Selected Answer: D

Even though replication is asynchronous I believe it is within the 100 ms requirement. It is also cost effective with autoscaling.

upvoted 1 times

 **KaranGandhi30** 3 years ago

Selected Answer: D

Option A? How can you use serverless and scale in the same statement

upvoted 1 times

🗳️ 👤 **VPup** 3 years, 3 months ago

Answer D

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Replication.html>

"As a result, all Aurora Replicas return the same data for query results with minimal replica lag. This lag is usually much less than 100 milliseconds after the primary instance has written an update. "

upvoted 3 times

🗳️ 👤 **Dip11** 3 years, 6 months ago

Option D makes most sense.

upvoted 2 times

🗳️ 👤 **LMax** 3 years, 6 months ago

Answer D

upvoted 4 times

🗳️ 👤 **myutran** 3 years, 7 months ago

Ans: D

upvoted 3 times

🗳️ 👤 **Robbb** 3 years, 7 months ago

A is a ridiculous answer. If you say A, don't bother taking the test.

upvoted 3 times

🗳️ 👤 **GeeBeeEl** 3 years, 7 months ago

Oh, I love this your response..... Option A seems to make the most sense to me and I will take the test

Here is a link to help with the correct response -- <https://aws.amazon.com/rds/aurora/faqs/>

Amazon Aurora Serverless is an on-demand, autoscaling configuration for the MySQL-compatible and PostgreSQL-compatible editions of Amazon Aurora. An Aurora Serverless DB cluster automatically starts up, shuts down, and scales capacity up or down based on your application's needs. Aurora Serverless provides a relatively simple, cost-effective option for infrequent, intermittent, or unpredictable workloads.

The question is focused on unpredictable workloads

upvoted 4 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

The answer is saying "Turn on Serverless option on the cluster". There's no such thing.

upvoted 1 times

🗳️ 👤 **[Removed]** 3 years, 6 months ago

you can't turn on serverless option.you need to take a snapshot and restore it to Aurora Serverless. A is definitely wrong

upvoted 5 times

🗳️ 👤 **Robbb** 3 years, 7 months ago

Also it says that the new workload is unpredictable, and yet should have no impact on the current operations. It takes time to adjust to unpredictable workloads, so D does not solve the stated problem.

upvoted 1 times

🗳️ 👤 **toppic26** 3 years, 6 months ago

it says read workload, not write.

upvoted 1 times

🗳️ 👤 **Robbb** 3 years, 7 months ago

B is the best choice. A cloned Cluster will use the existing DB cluster until those items are written over, so that will have the fastest, immediate response. The best solution, of course, is to dedicate a read replica to the team and use an instance endpoint. D does not directly address the issue.

upvoted 1 times

🗳️ 👤 **Dip11** 3 years, 6 months ago

Clone is a one time copy, it does not continuously replicate which is a requirement here.

upvoted 3 times



🗳️ 👤 **Glendon** 3 years, 7 months ago

Why can't B be an option? Create a clone of the Aurora cluster and use the clone for data read on the Dashboard?

upvoted 1 times

🗳️ 👤 **AM** 3 years, 7 months ago

B is not an option since clone does not support read scaling of the same cluster. This question is on read scaling of Aurora cluster. Only achieved with Read replicas. Serverless as pointed out is not available as an option. Correct an D.
upvoted 1 times

  **Faz** 3 years, 7 months ago

Ans is D.



<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Integrating.AutoScaling.html#:~:text=The%20scaling%20policy%20defines%20t>

upvoted 2 times

  **JobinAkaJoe** 3 years, 7 months ago

D is the best choice

upvoted 2 times

  **kilkar** 3 years, 7 months ago

Ashoks is right. There is no turn on option to Serverless, it needs migration

upvoted 3 times

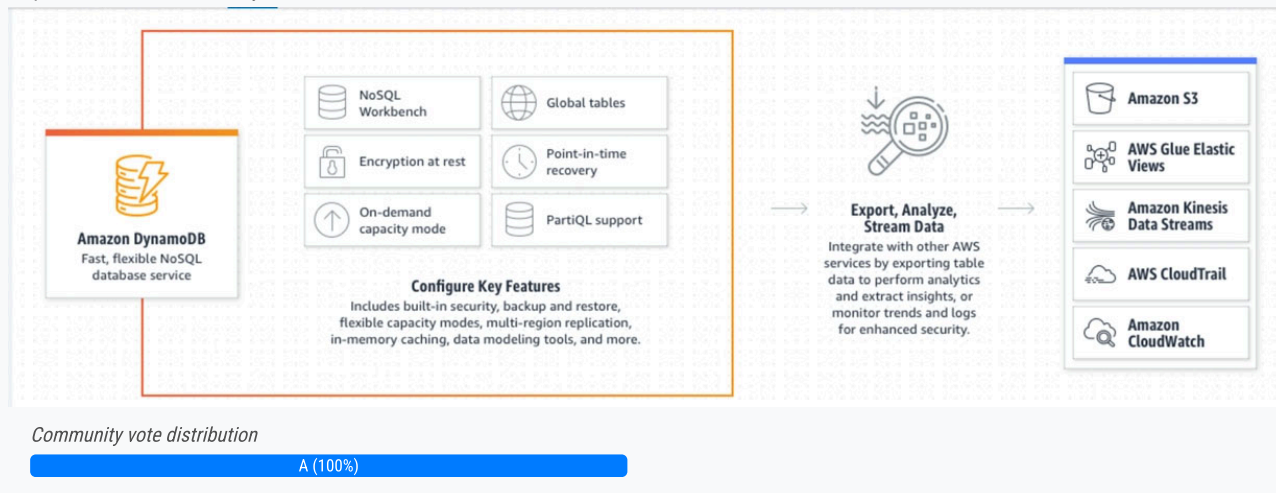
A retail company is about to migrate its online and mobile store to AWS. The company's CEO has strategic plans to grow the brand globally. A Database Specialist has been challenged to provide predictable read and write database performance with minimal operational overhead. What should the Database Specialist do to meet these requirements?

- A. Use Amazon DynamoDB global tables to synchronize transactions
- B. Use Amazon EMR to copy the orders table data across Regions
- C. Use Amazon Aurora Global Database to synchronize all transactions
- D. Use Amazon DynamoDB Streams to replicate all DynamoDB transactions and sync them

Suggested Answer: A

Reference:

<https://aws.amazon.com/dynamodb/>



awsmonster Highly Voted 3 years, 4 months ago

A

<https://aws.amazon.com/dynamodb/features/>

With global tables, your globally distributed applications can access data locally in the selected regions to get single-digit millisecond read and write performance.

Not Aurora Global Database, as per this link: https://aws.amazon.com/rds/aurora/global-database/?nc1=h_ls . Aurora Global Database lets you easily scale database reads across the world and place your applications close to your users.

upvoted 8 times

NishithShah Most Recent 1 year, 5 months ago

A. Use Amazon DynamoDB global tables to synchronize transactions

upvoted 1 times

nahunaws 1 year, 7 months ago

A is best choice <https://aws.amazon.com/dynamodb/features/>

upvoted 1 times

redman50 2 years, 2 months ago

Using Amazon DynamoDB global tables to synchronize transactions - is a viable option, but it's more suited for a NoSQL database. It can provide predictable read and write performance, but with DynamoDB, there are limitations regarding query and transactional capabilities.

upvoted 1 times

im_not_robot 2 years, 3 months ago

Selected Answer: A

A is correct

B is incorrect because aurora global db only allow 'write' on primary region

upvoted 4 times

🗨️ 👤 **novice_expert** 3 years ago

Selected Answer: A

A. Use Amazon DynamoDB global tables

upvoted 3 times

🗨️ 👤 **awsmonster** 3 years, 4 months ago

Ans: A

upvoted 2 times

🗨️ 👤 **jove** 3 years, 5 months ago

Selected Answer: A

I go with option A

upvoted 3 times

A company is closing one of its remote data centers. This site runs a 100 TB on-premises data warehouse solution. The company plans to use the AWS Schema Conversion Tool (AWS SCT) and AWS DMS for the migration to AWS. The site network bandwidth is 500 Mbps. A Database Specialist wants to migrate the on-premises data using Amazon S3 as the data lake and Amazon Redshift as the data warehouse. This move must take place during a 2-week period when source systems are shut down for maintenance. The data should stay encrypted at rest and in transit. Which approach has the least risk and the highest likelihood of a successful data transfer?

- A. Set up a VPN tunnel for encrypting data over the network from the data center to AWS. Leverage AWS SCT and apply the converted schema to Amazon Redshift. Once complete, start an AWS DMS task to move the data from the source to Amazon S3. Use AWS Glue to load the data from Amazon S3 to Amazon Redshift.
- B. Leverage AWS SCT and apply the converted schema to Amazon Redshift. Start an AWS DMS task with two AWS Snowball Edge devices to copy data from on-premises to Amazon S3 with AWS KMS encryption. Use AWS DMS to finish copying data to Amazon Redshift.
- C. Leverage AWS SCT and apply the converted schema to Amazon Redshift. Once complete, use a fleet of 10 TB dedicated encrypted drives using the AWS Import/Export feature to copy data from on-premises to Amazon S3 with AWS KMS encryption. Use AWS Glue to load the data to Amazon Redshift.
- D. Set up a VPN tunnel for encrypting data over the network from the data center to AWS. Leverage a native database export feature to export the data and compress the files. Use the `aws S3 cp` multi-port upload command to upload these files to Amazon S3 with AWS KMS encryption. Once complete, load the data to Amazon Redshift using AWS Glue.

Suggested Answer: C

Community vote distribution

B (100%)

  **pan24** Highly Voted 3 years, 8 months ago

Ans: B

<https://aws.amazon.com/blogs/database/new-aws-dms-and-aws-snowball-integration-enables-mass-database-migrations-and-migrations-of-large-databases/>

upvoted 19 times

  **BillyMadison** 3 years, 8 months ago

Going with B as well due to your compelling link.

upvoted 2 times

  **BillyMadison** 3 years, 7 months ago

Found this link that agrees with B

<https://docs.aws.amazon.com/SchemaConversionTool/latest/userguide/agents.dw.html>

"Large-scale data migrations can include many terabytes of information, and can be slowed by network performance and by the sheer amount of data that has to be moved. AWS Snowball Edge is an AWS service you can use to transfer data to the cloud at faster-than-network speeds using an AWS-owned appliance."

"When you use AWS SCT and an AWS Snowball Edge device, you migrate your data in two stages. First, you use the AWS SCT to process the data locally and then move that data to the AWS Snowball Edge device. You then send the device to AWS using the AWS Snowball Edge process, and then AWS automatically loads the data into an Amazon S3 bucket. Next, when the data is available on Amazon S3, you use AWS SCT to migrate the data to Amazon Redshift."

upvoted 4 times

  **szmulder** 3 years, 7 months ago

B is incorrect.

1. You just need one snowball edge device because one device can hold up to 80 TB of data.

2. The lastest step for B is "Use AWS DMS to finish copying data to Amazon Redshift." but from the AWS docs need "use AWS SCT to migrate the data to Amazon Redshift."

upvoted 2 times

  **BillyMadison** 3 years, 7 months ago

There is 100 TB of data that needs to be transferred. Since a snowball can only hold 80, it makes sense that there would be a need for 2 of them. Hence B.

upvoted 9 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

Wrong. you need TWO snowball edge devices since out of 100TB only 80TB is usable. This is a standard question in tutotrialsdojo as well - very respectable site.

upvoted 2 times

🗳️ 👤 **Exia** Highly Voted 🏆 3 years, 6 months ago

Although snowball is a petabyte-scale data transport solution, 100TB database migration through a 500 Mbps network is impossible in 2 weeks.

bandwidth: 500 Mbps = 62.5 MB/s

capacity: 100 TB = 100 * 10**6 = 100000000 MB

time consume: 100000000 / 62.5 MB/s = 1600000.0 s =

1600000 / 60 / 60 / 24 = 18.5 days

upvoted 5 times

🗳️ 👤 **58a2d17** 1 year, 1 month ago

That is why you need snowball. Network speed won't do it .

upvoted 1 times

🗳️ 👤 **jitesh_k** 1 year, 6 months ago

How is 500MBPS equal to 62.5 MB/sec?

upvoted 1 times

🗳️ 👤 **jitesh_k** 1 year, 6 months ago

Got it. Lowercase "b" is bits while uppercase "B" is bytes. 500 mbps is 500/8 MB/s.

upvoted 1 times

🗳️ 👤 **KaranGandhi30** Most Recent 🔍 3 years ago

Selected Answer: B

With the given bandwidth any other option will be completed in more than 20 days.

upvoted 4 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: B

<https://aws.amazon.com/blogs/database/new-aws-dms-and-aws-snowball-integration-enables-mass-database-migrations-and-migrations-of-large-databases/>

upvoted 1 times

🗳️ 👤 **jove** 3 years, 5 months ago

Option B.. With SCT data extraction agent you can extract the data as well as the schema. Snowball Edge devices are safe.

upvoted 1 times

🗳️ 👤 **LMax** 3 years, 6 months ago

Agree with B

upvoted 1 times

🗳️ 👤 **myutran** 3 years, 6 months ago

Ans: B

upvoted 1 times

🗳️ 👤 **JobinAkaJoe** 3 years, 6 months ago

B is the right choice here

upvoted 1 times

🗳️ 👤 **Ashoks** 3 years, 6 months ago

B. Snow ball will be the solution to transfer 100TB, since existing 500Mbps bandwidth is not enough

upvoted 2 times

🗳️ 👤 **AWSCert2020** 3 years, 7 months ago

For me B here

upvoted 1 times

🗳️ 👤 **szmulder** 3 years, 7 months ago

A. Answer: It's possible but copy 100TB via internet is not reliable

B. Answer: "Use AWS DMS to finish copying data to Amazon Redshift." but from the AWS docs need "use AWS SCT to migrate the data to Amazon Redshift."

C. Answer: Correct

D. Answer: It's possible but it's missing the step AWS SCT and apply the converted schema to Amazon Redshift.
upvoted 1 times

🗨️ 👤 **[Removed]** 3 years, 7 months ago
didn't know you can ship fleets of drives to AWS.
upvoted 2 times

🗨️ 👤 **szmulder** 3 years, 6 months ago
sorry, change my mind, B is correct.
upvoted 1 times

🗨️ 👤 **badrik** 3 years, 7 months ago
B is right option
upvoted 1 times

🗨️ 👤 **BillyC** 3 years, 7 months ago
B is correct for me
upvoted 1 times

A company is looking to migrate a 1 TB Oracle database from on-premises to an Amazon Aurora PostgreSQL DB cluster. The company's Database Specialist discovered that the Oracle database is storing 100 GB of large binary objects (LOBs) across multiple tables. The Oracle database has a maximum LOB size of 500 MB with an average LOB size of 350 MB. The Database Specialist has chosen AWS DMS to migrate the data with the largest replication instances.

How should the Database Specialist optimize the database migration using AWS DMS?

- A. Create a single task using full LOB mode with a LOB chunk size of 500 MB to migrate the data and LOBs together
- B. Create two tasks: task1 with LOB tables using full LOB mode with a LOB chunk size of 500 MB and task2 without LOBs
- C. Create two tasks: task1 with LOB tables using limited LOB mode with a maximum LOB size of 500 MB and task 2 without LOBs
- D. Create a single task using limited LOB mode with a maximum LOB size of 500 MB to migrate data and LOBs together

Suggested Answer: C

Community vote distribution

B (67%)

C (33%)

  **vicks316** Highly Voted 3 years, 7 months ago

C sounds correct since as per link https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html#CHAP_BestPractices.LOBS, "AWS DMS migrates LOB data in two phases:

1. AWS DMS creates a new row in the target table and populates the row with all data except the associated LOB value.
2. AWS DMS updates the row in the target table with the LOB data."

This means that we would need two tasks, one per phase and use limited LOB mode for best performance.

upvoted 17 times

  **khchan123** 3 years ago

I think it is C. The explanation from vicks316 is irrelevant. From this link ->

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Tasks.LOBSupport.html

Using the Max LOB size (K) option with a value greater than 63KB impacts the performance of a full load configured to run in limited LOB mode.

During a full load, DMS allocates memory by multiplying the Max LOB size (k) value by the Commit rate, and the product is multiplied by the number of LOB columns. When DMS can't pre-allocate that memory, DMS starts consuming SWAP memory, and that impacts performance of a full load. So, if you experience performance issues when using limited LOB mode, consider decreasing the commit rate until you achieve an acceptable level of performance. You can also consider using inline LOB mode for supported endpoints once you understand your LOB distribution for the table.

upvoted 4 times

  **JobinAkaJoe** 3 years, 7 months ago

This doesn't mean we need two tasks. It just explains how LOBs are copied to target in any scenario. First all columns except LOBs then the LOB.

upvoted 2 times

  **Shunpin** Highly Voted 3 years, 5 months ago

Selected Answer: B

Limited LOB maxsize is 100MB. This question is about LOB average 350MB. Full LOB is better option

upvoted 14 times

  **Sathish_dbs** 1 year, 7 months ago

the link says it not a hard limit 100MB, it is recommended.

upvoted 1 times

  **Shunpin** 3 years, 5 months ago

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Tasks.LOBSupport.html

upvoted 4 times

  **MultiAZ** Most Recent 1 year, 4 months ago

Selected Answer: C

I believe the right answer is C

As for B, 500 chunk size is way too big and will lead to many failures.

upvoted 1 times

🗨️ **rrshah83** 1 year, 5 months ago

Selected Answer: C

Limited LOB with max size of 500MB makes migration fast.

upvoted 1 times

🗨️ **jitesh_k** 1 year, 6 months ago

Full LOB mode has no information about chunk size.

Full LOB mode – In full LOB mode AWS DMS migrates all LOBs from source to target regardless of size. In this configuration, AWS DMS has no information about the maximum size of LOBs to expect. Thus, LOBs are migrated one at a time, piece by piece. Full LOB mode can be quite slow.

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Tasks.LOBSupport.html

Either C or D - but C seems more optimal.

upvoted 1 times

🗨️ **Sathish_dbs** 1 year, 7 months ago

the question is confusing, but the lesson here full lob -> incur performance issue, limited lob - smooth, hopefully the same question doesn't repeat, if so then let it come with some additional clarity on performance, either way 500MB is huge and we expect performance issue, this seems to be meaningless question

upvoted 1 times

🗨️ **roymunson** 1 year, 8 months ago

Selected Answer: B

It can not be C because of

"Limited LOB mode – In limited LOB mode, you set a maximum LOB size for DMS to accept. That enables DMS to pre-allocate memory and load the LOB data in bulk. LOBs that exceed the maximum LOB size are truncated, and a warning is issued to the log file. In limited LOB mode, you can gain significant performance over full LOB mode. We recommend that you use limited LOB mode whenever possible. The maximum recommended value is 102400 KB (100 MB)."

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Tasks.LOBSupport.html

upvoted 2 times

🗨️ **cuibap** 1 year, 10 months ago

To optimize for large size LOB, we should use multiple tasks => eliminate A & D.

B: we can't configure such large LOB chunk size of 500 MB. Please note that "chunk size" is different from "max LOB size".

So I select C but can anyone explain what benefit of task 2 without LOBs?

upvoted 1 times

🗨️ **rn30** 1 year, 10 months ago

B looks to be right

Task1 with LOB Tables:

Use the full LOB mode in AWS DMS to ensure efficient streaming of LOB data during replication.

Set the LOB chunk size to match the maximum LOB size in Oracle (500 MB). This ensures optimal data transfer and minimizes network overhead.

Task2 without LOBs:

Migrate the remaining tables that do not contain LOBs separately in a task without LOB-specific configurations.

upvoted 2 times

🗨️ **saikarthikeya777** 1 year, 11 months ago

Selected Answer: C

full lob is quite slow ! to optimize we need to use limited lob even if 500mb is not recommended size . Maybe i am super confused

upvoted 2 times

🗨️ **backbencher2022** 2 years, 2 months ago

B is correct. Option C could have been correct however, please note the word "recommended size" for Limited LOB mode which is 100MB. Agreed that 100 MB is not the max limit for limited LOB however, anything above 100MB in limited LOB is not recommended which makes full LOB with 2 different tasks a better option.

upvoted 2 times

🗨️ **anantarb** 2 years, 4 months ago

Selected Answer: B

The Max recommended LOB size is 100MB in limited LOB mode. So can't be C. B is correct.

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Tasks.LOBSupport.html

upvoted 3 times

🗨️ 👤 **teo2157** 2 years, 4 months ago

Selected Answer: C

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html#CHAP_BestPractices.LOBS. Limited LOB mode migrates all LOB values up to a user-specified size limit (default is 32 KB). LOB values larger than the size limit must be manually migrated. Limited LOB mode, the default for all migration tasks, typically provides the best performance. However, ensure that the Max LOB size parameter setting is correct. Set this parameter to the largest LOB size for all your tables.

upvoted 1 times

🗨️ 👤 **parle101** 2 years, 4 months ago

Selected Answer: C

C is the answer - <https://aws.amazon.com/premiumsupport/knowledge-center/dms-improve-speed-lob-data/>

upvoted 1 times

🗨️ 👤 **khun** 2 years, 5 months ago

Selected Answer: C

C. is the answer. Faster = 2 task Limited Lob and No Lob

Full LOB mode – Migrate complete LOBs regardless of size. AWS DMS migrates LOBs piecewise in chunks controlled by the Max LOB Size parameter. This mode is slower than using limited LOB mode.

Limited LOB mode – Truncate LOBs to the value specified by the Max LOB Size parameter. This mode is faster than using full LOB mode.

upvoted 2 times

🗨️ 👤 **shuo82** 2 years, 5 months ago

Selected Answer: C

C is correct, 100MB is the maximum recommended value, not a limitation.

B is wrong, because 500MB is too huge for chunk setting, connection error would be encountered.

When a task is configured to use full LOB mode, AWS DMS retrieves LOBs in pieces. The LOB chunk size (K) option determines the size of each piece. When setting this option, pay particular attention to the maximum packet size allowed by your network configuration. If the LOB chunk size exceeds your maximum allowed packet size, you might see disconnect errors. The recommended value for LobChunkSize is 64 kilobytes. Increasing the value for LobChunkSize above 64 kilobytes can cause task failures.

upvoted 2 times

🗨️ 👤 **rag1482** 2 years, 7 months ago

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html#CHAP_BestPractices.LOBS

Answer : C

upvoted 1 times

A Database Specialist is designing a disaster recovery strategy for a production Amazon DynamoDB table. The table uses provisioned read/write capacity mode, global secondary indexes, and time to live (TTL). The Database Specialist has restored the latest backup to a new table. To prepare the new table with identical settings, which steps should be performed? (Choose two.)

- A. Re-create global secondary indexes in the new table
- B. Define IAM policies for access to the new table
- C. Define the TTL settings
- D. Encrypt the table from the AWS Management Console or use the update-table command
- E. Set the provisioned read and write capacity

Suggested Answer: AE

Reference:

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

Read/Write Capacity Mode

PDF | Kindle | RSS

Amazon DynamoDB has two read/write capacity modes for processing reads and writes on your tables:

- On-demand
- Provisioned (default, free-tier eligible)

The read/write capacity mode controls how you are charged for read and write throughput and how you manage capacity. You can set the read/write capacity mode when creating a table or you can change it later.

Community vote distribution

BC (100%)

 **Nots**  3 years, 7 months ago

I'll go B and C.

The following items need to be reconfigured after restoring the DynamoDB table.

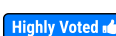
- AutoScaling policy
- IAM policy
- CloudWatch settings
- Tags
- Stream settings
- TTL

upvoted 16 times

 **johnconnor** 3 years, 7 months ago

You are right, it is explained here -> https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/backuprestore_HowItWorks.html

upvoted 3 times


 **RotterDam**  3 years, 2 months ago

Selected Answer: BC

The following are not restored and would need to be configured again:

- IAM and Autoscaling Policies
- Cloudwatch Triggers and Alarms
- TTL and Streams
- Tag

upvoted 6 times

 **ken_test1234**  2 years, 2 months ago

Selected Answer: BC

The answer is BC because this is what included in backup

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

upvoted 4 times

  **novice_expert** 3 years ago

Selected Answer: BC

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

You must manually set up the following on the restored table:

Auto scaling policies

AWS Identity and Access Management (IAM) policies

Amazon CloudWatch metrics and alarms

Tags

Stream settings

Time to Live (TTL) settings

upvoted 5 times

  **GMartinelli** 3 years, 5 months ago

Selected Answer: BC

Option B, C

upvoted 3 times

A Database Specialist is creating Amazon DynamoDB tables, Amazon CloudWatch alarms, and associated infrastructure for an Application team using a development AWS account. The team wants a deployment method that will standardize the core solution components while managing environment-specific settings separately, and wants to minimize rework due to configuration errors.

Which process should the Database Specialist recommend to meet these requirements?

- A. Organize common and environmental-specific parameters hierarchically in the AWS Systems Manager Parameter Store, then reference the parameters dynamically from an AWS CloudFormation template. Deploy the CloudFormation stack using the environment name as a parameter.
- B. Create a parameterized AWS CloudFormation template that builds the required objects. Keep separate environment parameter files in separate Amazon S3 buckets. Provide an AWS CLI command that deploys the CloudFormation stack directly referencing the appropriate parameter bucket.
- C. Create a parameterized AWS CloudFormation template that builds the required objects. Import the template into the CloudFormation interface in the AWS Management Console. Make the required changes to the parameters and deploy the CloudFormation stack.
- D. Create an AWS Lambda function that builds the required objects using an AWS SDK. Set the required parameter values in a test event in the Lambda console for each environment that the Application team can modify, as needed. Deploy the infrastructure by triggering the test event in the console.

Suggested Answer: C

Reference:

<https://aws.amazon.com/blogs/mt/aws-cloudformation-signed-sealed-and-deployed/>

Why State Street chose AWS CloudFormation

State Street considered tools such as Terraform, Puppet, and Chef to help enable Infrastructure as Code for its environments. **AWS CloudFormation** was selected as a preferred platform, due to key differentiators:

1. Integrated resource management for AWS services:

- 1. A singular model for resource lifecycle management across the various services.
- 2. An advantage over **AWS Serverless Application Model (SAM)** and **AWS CDK** (at this time), as these are more specific to a given service.
- 3. Users manage resources via the console, programmatically, and various CI/CD tools. CloudFormation provides oversight of concurrency and state management for more effective management.

2. Bimodal application deployment:

- 1. Deployment for applications within an AWS account (and potentially a Region).
- 2. Deployment of complex application across multiple AWS accounts and multiple Regions (through the use of **AWS CloudFormation StackSets**).
- 3. A programmatic interface for **configuration drift detection**, which can be used to trigger compensating actions, alerts or other actions.

Community vote distribution

A (100%)

 **chicagomassageseeker** Highly Voted 3 years, 8 months ago

A. AWS Systems Manager Parameter Store
upvoted 10 times

 **BillyMadison** 3 years, 7 months ago

Is it A based off this?

<https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-parameter-store.html>

upvoted 3 times

 **pan24** 3 years, 7 months ago

A is correct

<https://aws.amazon.com/blogs/mt/integrating-aws-cloudformation-with-aws-systems-manager-parameter-store/>

upvoted 5 times

 **kyyo** Most Recent 1 year, 3 months ago

C is also possible. However, A is better than C.

upvoted 1 times

🗨️ 👤 **SuriSagar** 1 year, 7 months ago

The answer is A

upvoted 1 times

🗨️ 👤 **lollyj** 2 years, 5 months ago

Selected Answer: A

Parameter store to separate the environments

upvoted 3 times

🗨️ 👤 **sachin** 2 years, 10 months ago

Mapping should have been the right approach for handing environment and it related setting but that option is not there..

We can still achieve setting as using parameters.

A

upvoted 1 times

🗨️ 👤 **novice_expert** 3 years ago

Selected Answer: A

A. Organize common and environmental-specific parameters hierarchically in the AWS Systems Manager Parameter Store, then reference the parameters dynamically from an AWS CloudFormation template. Deploy the CloudFormation stack using the environment name as a parameter.

upvoted 3 times

🗨️ 👤 **LMax** 3 years, 6 months ago

The only reasonable answer is A

upvoted 2 times

🗨️ 👤 **myutran** 3 years, 6 months ago

Ans: A

upvoted 1 times

🗨️ 👤 **JobinAkaJoe** 3 years, 6 months ago

Not sure. I will go with A

upvoted 1 times

🗨️ 👤 **Ashoks** 3 years, 7 months ago

I would select A

upvoted 2 times

🗨️ 👤 **Manmohan** 3 years, 7 months ago

I will go with B

upvoted 2 times

🗨️ 👤 **AWSCert2020** 3 years, 7 months ago

For me is C or B because on DynamoDB probably i need different configuration (for example RCU and WCU) between Dev and Prod environment.

Probably i choose C because CLI not support parameters on S3 but only the body template.

upvoted 1 times

🗨️ 👤 **BillyC** 3 years, 7 months ago

A is correct

upvoted 2 times

A company runs online transaction processing (OLTP) workloads on an Amazon RDS for PostgreSQL Multi-AZ DB instance. Tests were run on the database after work hours, which generated additional database logs. The free storage of the RDS DB instance is low due to these additional logs. What should the company do to address this space constraint issue?

- A. Log in to the host and run the `rm $PGDATA/pg_logs/*` command
- B. Modify the `rds.log_retention_period` parameter to 1440 and wait up to 24 hours for database logs to be deleted
- C. Create a ticket with AWS Support to have the logs deleted
- D. Run the `SELECT rds_rotate_error_log()` stored procedure to rotate the logs

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **BillyMadison** Highly Voted 3 years, 8 months ago

I cant find anything for the other answers, so I'm going with B based off

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_LogAccess.Concepts.PostgreSQL.html

"To set the retention period for system logs, use the `rds.log_retention_period` parameter. You can find `rds.log_retention_period` in the DB parameter group associated with your DB instance. The unit for this parameter is minutes. For example, a setting of 1,440 retains logs for one day. The default value is 4,320 (three days)."

"If storage gets too low, Aurora might delete compressed PostgreSQL logs before the retention period expires. If logs are deleted early, you get a message like the following.

The oldest PostgreSQL log files were deleted due to local storage constraints."

upvoted 7 times

🗳️ 👤 **trietnv** 2 years, 11 months ago

Reason: DB logs (error files) that are retained for too long.

Because by default, PostgreSQL error log files have a retention value of 4,320 minutes (three days). Large log files can use more space because of higher workloads. You can change the retention period for system logs using the `rds.log_retention_period` parameter in the DB parameter group associated with your DB instance. For example, if you set the value to 1440, then logs are retained for one day

Ref: <https://aws.amazon.com/premiumsupport/knowledge-center/diskfull-error-rds-postgresql/>

upvoted 1 times

🗳️ 👤 **adelcold** Most Recent 1 year, 11 months ago

Selected Answer: B

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_LogAccess.Concepts.PostgreSQL.html

upvoted 1 times

🗳️ 👤 **redman50** 2 years, 2 months ago

In this scenario, the company is facing a low storage issue due to additional logs. To address this issue, they can rotate the logs to free up storage space using the stored procedure `rds_rotate_error_log()`. This procedure rotates the current log file and starts a new one. The rotated logs are compressed and stored in the log directory, freeing up space in the main storage.

Option A is incorrect because removing log files manually is not recommended, as it may cause issues and loss of data.

Option B is incorrect because changing the log retention period will not delete the existing logs immediately, and the company needs to wait for up to 24 hours for the logs to be deleted.

Option C is incorrect because AWS support does not provide log deletion services.

upvoted 3 times

🗳️ 👤 **jitesh_k** 1 year, 6 months ago

There is no function named `rds_rotate_error_log()`. It is just a distractor. Check this link:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_LogAccess.Concepts.PostgreSQL.html

upvoted 1 times

🗳️ 👤 **SachinGoel** 2 years, 7 months ago

Selected Answer: B

B is correct



upvoted 1 times

  **novice_expert** 3 years ago

Selected Answer: B

reduce rds.log_retention_period parameter and wait



upvoted 2 times

  **tugboat** 3 years, 3 months ago

Selected Answer: B

correct



upvoted 1 times

  **AmitB** 3 years, 4 months ago

B is correct Answer.

To set the retention period for system logs, use the `rds.log_retention_period` parameter. You can find `rds.log_retention_period` in the DB parameter group associated with your DB instance. The unit for this parameter is minutes. For example, a setting of 1,440 retains logs for one day. The default value is 4,320 (three days). The maximum value is 10,080 (seven days).

upvoted 3 times



  **guru_ji** 3 years, 6 months ago

B is Correct.

==> Anyone planning for exam?

We can share study material with each other, it would be beneficial for both. You can email me on "awsdbguru at gmail"

upvoted 3 times


  **guru_ji** 3 years, 6 months ago

Correct Answer is ==>> B

any idea how much Q we will get in real exam from Q available here?

anyone is preparing for this exam and want to do group study with us, comment with mail_id.

upvoted 2 times

  **LMax** 3 years, 7 months ago

My choice is B


upvoted 2 times

  **AWSCert2020** 3 years, 7 months ago

B here

The `SELECT rds_rotate_error_log()` does not exist on RDS

upvoted 2 times

  **BillyC** 3 years, 7 months ago

B is correct

upvoted 3 times

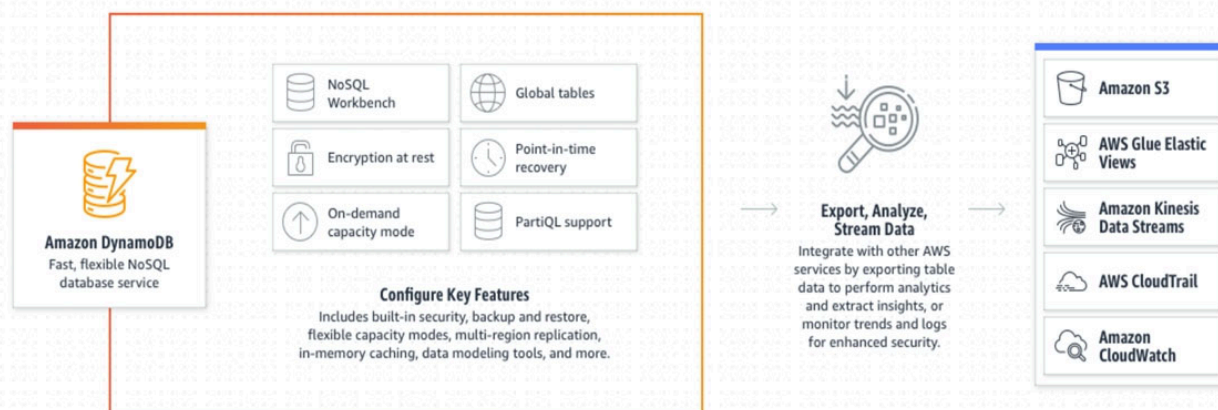
A user has a non-relational key-value database. The user is looking for a fully managed AWS service that will offload the administrative burdens of operating and scaling distributed databases. The solution must be cost-effective and able to handle unpredictable application traffic. What should a Database Specialist recommend for this user?

- A. Create an Amazon DynamoDB table with provisioned capacity mode
- B. Create an Amazon DocumentDB cluster
- C. Create an Amazon DynamoDB table with on-demand capacity mode
- D. Create an Amazon Aurora Serverless DB cluster

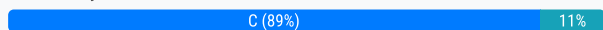
Suggested Answer: C

Reference:

<https://aws.amazon.com/dynamodb/>



Community vote distribution



jove Highly Voted 3 years, 5 months ago

Selected Answer: C

Key-value database -> DynamoDB

Capable of dealing with unexpected application traffic -> on-demand capacity mode

upvoted 6 times

Pranava_GCP Most Recent 1 year, 8 months ago

Selected Answer: C

C. Create an Amazon DynamoDB table with on-demand capacity mode

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.html#HowItWorks.OnDemand>
upvoted 1 times

aviathor 1 year, 12 months ago

Selected Answer: A

"With provisioned capacity you can also use auto scaling to automatically adjust your table's capacity based on the specified utilization rate to ensure application performance, and also to potentially reduce costs. To configure auto scaling in DynamoDB, set the minimum and maximum levels of read and write capacity in addition to the target utilization percentage."

<https://docs.aws.amazon.com/wellarchitected/latest/serverless-applications-lens/capacity.html>
upvoted 1 times

ajndcjeandca 1 year, 9 months ago

If auto scaling was mentioned in the option then that would've been the most cost effective solution but since it is not mentioned C seems to be more appropriate
upvoted 1 times

🗨️ 👤 **novice_expert** 3 years ago

Selected Answer: C

C. Create an Amazon DynamoDB table with on-demand capacity mode
upvoted 1 times

🗨️ 👤 **2025flakyt** 3 years, 5 months ago

C is the answer

A key-value database is a type of nonrelational database that uses a simple key-value method to store data. A key-value database stores data as a collection of key-value pairs in which a key serves as a unique identifier.

On-demand mode is a good option to create new tables with unknown workloads.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.html#HowItWorks.OnDemand>

upvoted 2 times

A gaming company is designing a mobile gaming app that will be accessed by many users across the globe. The company wants to have replication and full support for multi-master writes. The company also wants to ensure low latency and consistent performance for app users. Which solution meets these requirements?

- A. Use Amazon DynamoDB global tables for storage and enable DynamoDB automatic scaling
- B. Use Amazon Aurora for storage and enable cross-Region Aurora Replicas
- C. Use Amazon Aurora for storage and cache the user content with Amazon ElastiCache
- D. Use Amazon Neptune for storage

Suggested Answer: A

Reference:

<https://aws.amazon.com/blogs/database/how-to-use-amazon-dynamodb-global-tables-to-power-multiregion-architectures/>

How to use Amazon DynamoDB global tables to power multiregion architectures

by Adrian Hornsby | on 20 DEC 2018 | in [Amazon DynamoDB, Database](#) | [Permalink](#) | [Comments](#) | [Share](#)

More and more, AWS customers want to make their applications available to globally dispersed users by deploying their application in multiple AWS Regions. These global users expect fast application performance.

In this post, I describe how to use [Amazon DynamoDB](#) to power the database of a global backend deployed in multiple AWS Regions. I use DynamoDB [global tables](#), which provide a fully managed, multiregion, and multimaster database so that you can deliver low-latency data access to your users no matter where they are located on the globe.

Why use a multiregion architecture?

AWS customers typically want a multiregion architecture for two reasons:

Community vote distribution

A (100%)

🗲️ 👤 **Pranava_GCP** 1 year, 8 months ago

Selected Answer: A

A. Use Amazon DynamoDB global tables for storage and enable DynamoDB automatic scaling
upvoted 2 times

🗲️ 👤 **Pranava_GCP** 1 year, 8 months ago

<https://aws.amazon.com/dynamodb/global-tables/>

upvoted 1 times

🗲️ 👤 **novice_expert** 3 years ago

Selected Answer: A

DynamoDB global tables have read and write in all regions, so these are kind of multi master like Aurora Multi Master

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/V2globaltables_HowItWorks.html

upvoted 3 times

🗲️ 👤 **jove** 3 years, 5 months ago

Option A

upvoted 1 times

🗲️ 👤 **GMartinelli** 3 years, 6 months ago

Selected Answer: A

Option A. I was looking for Aurora Multi-Master or DynamoDB Global Tables, there is only one of them.

upvoted 1 times

🗲️ 👤 **leunamE** 3 years, 6 months ago

Option A. worldwide, multi-master writes, minimal latency
upvoted 1 times

A Database Specialist needs to speed up any failover that might occur on an Amazon Aurora PostgreSQL DB cluster. The Aurora DB cluster currently includes the primary instance and three Aurora Replicas.

How can the Database Specialist ensure that failovers occur with the least amount of downtime for the application?

- A. Set the TCP keepalive parameters low
- B. Call the AWS CLI failover-db-cluster command
- C. Enable Enhanced Monitoring on the DB cluster
- D. Start a database activity stream on the DB cluster

Suggested Answer: B

Community vote distribution

A (100%)

🗳️ 👤 **pan24** Highly Voted 👍 3 years, 8 months ago

Ans: A

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.BestPractices.html#AuroraPostgreSQL.BestPractices.FastFailover.T>
upvoted 13 times

🗳️ 👤 **gelsm** 3 years, 7 months ago

A

"Enabling TCP keepalive parameters and setting them aggressively ensures that if your client is no longer able to connect to the database, then any active are quickly closed. This action allows the application to react appropriately, such as by picking a new host to connect to."

upvoted 3 times

🗳️ 👤 **guru_ji** 3 years, 7 months ago

Correct Answer: A

upvoted 2 times

🗳️ 👤 **adelcold** Most Recent 🕒 1 year, 11 months ago

Selected Answer: A

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.BestPractices.FastFailover.html>
upvoted 3 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: A

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.BestPractices.html#AuroraPostgreSQL.BestPractices.FastFailover.T>

You need to set the following TCP keepalive parameters:

tcp_keepalive_time controls the time, in seconds, after which a keepalive packet is sent when no data has been sent by the socket (ACKs are not considered recommend the following setting:

tcp_keepalive_time = 1

tcp_keepalive_intvl controls the time, in seconds, between sending subsequent keepalive packets after the initial packet is sent (set using the tcp_keepalive_ parameter). We recommend the following setting:

tcp_keepalive_intvl = 1

tcp_keepalive_probes is the number of unacknowledged keepalive probes that occur before the application is notified. We recommend the following setting:

tcp_keepalive_probes = 5

These settings should notify the application within five seconds when the database stops responding.

upvoted 4 times

🗨️ 👤 **guru_ji** 3 years, 7 months ago

Correct Answer: A

upvoted 1 times

🗨️ 👤 **LMax** 3 years, 7 months ago

Answer A

upvoted 2 times

🗨️ 👤 **myutran** 3 years, 7 months ago

Ans: A

upvoted 1 times

🗨️ 👤 **JobinAkaJoe** 3 years, 7 months ago

Option A

upvoted 1 times

🗨️ 👤 **Ashoks** 3 years, 7 months ago

Yes. It is A

upvoted 2 times

🗨️ 👤 **BillyC** 3 years, 8 months ago

A is correct

upvoted 1 times

🗨️ 👤 **Mickysingh** 3 years, 8 months ago

Ans A because we can reduce it by keeping low value of TCP among other parameter.

upvoted 2 times

🗨️ 👤 **chicagomassageseeker** 3 years, 8 months ago

Answer A

upvoted 1 times

A Database Specialist needs to define a database migration strategy to migrate an on-premises Oracle database to an Amazon Aurora MySQL DB cluster. The company requires near-zero downtime for the data migration. The solution must also be cost-effective. Which approach should the Database Specialist take?

- A. Dump all the tables from the Oracle database into an Amazon S3 bucket using datapump (expdp). Run data transformations in AWS Glue. Load the data from the S3 bucket to the Aurora DB cluster.
- B. Order an AWS Snowball appliance and copy the Oracle backup to the Snowball appliance. Once the Snowball data is delivered to Amazon S3, create a new Aurora DB cluster. Enable the S3 integration to migrate the data directly from Amazon S3 to Amazon RDS.
- C. Use the AWS Schema Conversion Tool (AWS SCT) to help rewrite database objects to MySQL during the schema migration. Use AWS DMS to perform the full load and change data capture (CDC) tasks.
- D. Use AWS Server Migration Service (AWS SMS) to import the Oracle virtual machine image as an Amazon EC2 instance. Use the Oracle Logical Dump utility to migrate the Oracle data from Amazon EC2 to an Aurora DB cluster.

Suggested Answer: D

Community vote distribution

C (100%)

  **chicagomassageseeker** Highly Voted 3 years, 8 months ago

Answer C. This is heterogenous migration and requires DMS and SCT.
upvoted 17 times

  **learnaws** Highly Voted 3 years, 7 months ago


Yep, C here
upvoted 6 times

  **BillyMadison** 3 years, 7 months ago

Looks correct
<https://aws.amazon.com/blogs/database/migrating-oracle-databases-with-near-zero-downtime-using-aws-dms/>
upvoted 2 times

  **renfdo** Most Recent 2 years, 3 months ago

Selected Answer: C
Answer C. This is heterogenous migration and requires DMS and SCT.
upvoted 1 times

  **examineme** 2 years, 5 months ago

Selected Answer: C
Answer is C: Use AWS Schema Conversion Tool to convert Database and DMS for fast transfer.
upvoted 2 times

  **rag1482** 2 years, 7 months ago

Answer: C
<https://docs.aws.amazon.com/dms/latest/sbs/chap-on-premoracle2aurora.html>
upvoted 2 times

  **novice_expert** 3 years ago

Selected Answer: C
<https://aws.amazon.com/blogs/database/migrating-oracle-databases-with-near-zero-downtime-using-aws-dms/>

Zero down time = First SCT and then DMS full and CDC

C. Use the AWS Schema Conversion Tool (AWS SCT) to help rewrite database objects to MySQL during the schema migration. Use AWS DMS to perform the full load and change data capture (CDC) tasks.
upvoted 3 times

  **LMax** 3 years, 7 months ago

Answer C should be the one.

upvoted 2 times

🗨️ 👤 **myutran** 3 years, 7 months ago

Ans: C

upvoted 1 times

🗨️ 👤 **JobinAkaJoe** 3 years, 7 months ago

C is the right choice

upvoted 1 times

🗨️ 👤 **Ashoks** 3 years, 7 months ago

Yes, it is C. Zero down time = First SCT and then DMS full and CDC

upvoted 2 times

🗨️ 👤 **Mickysingh** 3 years, 8 months ago

Ans D can not be right as it is not doing schema conversion and using oracle backup for restoring. Also downtime is high.

upvoted 2 times

A marketing company is using Amazon DocumentDB and requires that database audit logs be enabled. A Database Specialist needs to configure monitoring so that all data definition language (DDL) statements performed are visible to the Administrator. The Database Specialist has set the `audit_logs` parameter to enabled in the cluster parameter group.

What should the Database Specialist do to automatically collect the database logs for the Administrator?

- A. Enable DocumentDB to export the logs to Amazon CloudWatch Logs
- B. Enable DocumentDB to export the logs to AWS CloudTrail
- C. Enable DocumentDB Events to export the logs to Amazon CloudWatch Logs
- D. Configure an AWS Lambda function to download the logs using the `download-db-log-file-portion` operation and store the logs in Amazon S3

Suggested Answer: A

Reference:

<https://docs.aws.amazon.com/documentdb/latest/developerguide/profiling.html>

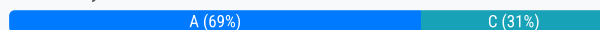
Profiling Amazon DocumentDB Operations

PDF | Kindle | RSS

You can use the profiler in Amazon DocumentDB (with MongoDB compatibility) to log the execution time and details of operations that were performed on your cluster. The profiler is useful for monitoring the slowest operations on your cluster to help you improve individual query performance and overall cluster performance.

By default, the profiler feature is disabled. When enabled, the profiler logs operations that are taking longer than a customer-defined threshold value (for example, 100 ms) to Amazon CloudWatch Logs. Logged details include the profiled command, time, plan summary,

Community vote distribution



leuname Highly Voted 3 years, 6 months ago

Option A.

upvoted 13 times

jeyp12 Highly Voted 3 years, 3 months ago

Looks like option C was added as a distractor. When the value of the `audit_logs` cluster parameter is enabled, you must also enable Amazon DocumentDB to export logs to Amazon CloudWatch. If you omit either of these steps, audit logs will not be sent to CloudWatch. Answer should be A

upvoted 13 times

MultiAZ Most Recent 1 year, 4 months ago

Selected Answer: A

Option A

- enable audit log

- enable export to CloudWatch Logs

upvoted 1 times

Pranava_GCP 1 year, 8 months ago

Selected Answer: A

A. Enable DocumentDB to export the logs to Amazon CloudWatch Logs

<https://docs.aws.amazon.com/documentdb/latest/developerguide/event-auditing.html>

" When auditing is enabled, Amazon DocumentDB exports your cluster's auditing records (JSON documents) to Amazon CloudWatch Logs. You can use Amazon CloudWatch Logs to analyze, monitor, and archive your Amazon DocumentDB auditing events."

upvoted 1 times

🗳️ 👤 **IhorK** 1 year, 9 months ago

Selected Answer: A

Amazon DocumentDB auditing supports the following event categories:

- Data Definition Language (DDL)
- Data Manipulation Language(DML)

Enabling auditing on a cluster is a two-step process.

Step 1. Enable the audit_logs cluster parameter (done)

Step 2. Enable Amazon CloudWatch Logs Export

<https://docs.aws.amazon.com/documentdb/latest/developerguide/event-auditing.html>

upvoted 1 times

🗳️ 👤 **adelcold** 1 year, 11 months ago

Selected Answer: A

<https://docs.aws.amazon.com/documentdb/latest/developerguide/event-auditing.html>

upvoted 1 times

🗳️ 👤 **f___16** 2 years ago

The audit logs are for Document DB events.

But to enable it, we must enable CloudWatch Logs exports on Document DB not Document DB Events.

upvoted 2 times

🗳️ 👤 **Pankaj24hrs** 2 years ago

Option A

Below is from AWS Documentation -

<https://docs.aws.amazon.com/documentdb/latest/developerguide/event-auditing.html>

When the value of the audit_logs cluster parameter is enabled, you must also enable Amazon DocumentDB to export logs to Amazon CloudWatch. If you omit either of these steps, audit logs will not be sent to CloudWatch.

upvoted 2 times

🗳️ 👤 **backbencher2022** 2 years, 2 months ago

Selected Answer: A

A is the correct option as this is a mandatory step besides changing the value of audit_logs parameter

upvoted 1 times

🗳️ 👤 **backbencher2022** 2 years, 2 months ago

A is the correct answer

upvoted 1 times

🗳️ 👤 **Nice_Guy** 2 years, 2 months ago

I don't know if this sounds like a silly question, but how about B?

upvoted 1 times

🗳️ 👤 **renfdo** 2 years, 3 months ago

Selected Answer: A

Answer is A with sure. When you go to the console and hit on create a cluster, you will see an option to enable exports Profile and Audit logs to cloudwatch.

upvoted 1 times

🗳️ 👤 **Arun32** 2 years, 6 months ago

Option A is correct. Just validated from AWS site:

When the value of the audit_logs cluster parameter is enabled, you must also enable Amazon DocumentDB to export logs to Amazon CloudWatch. If you omit either of these steps, audit logs will not be sent to CloudWatch.

upvoted 2 times

🗳️ 👤 **awsjji** 2 years, 7 months ago

Selected Answer: A

I am leaning towards A.

upvoted 1 times

🗨️ 👤 **Chirantan** 2 years, 10 months ago

I think Answer is A

When auditing is enabled, Amazon DocumentDB exports your cluster's auditing records (JSON documents) to Amazon CloudWatch Logs. You can use Amazon CloudWatch Logs to analyze, monitor, and archive your Amazon DocumentDB auditing events.

upvoted 3 times

🗨️ 👤 **kush_sumit** 2 years, 10 months ago

Selected Answer: A

Answer A:

<https://docs.aws.amazon.com/documentdb/latest/developerguide/event-auditing.html>

Step 2. Enable Amazon CloudWatch Logs Export

When the value of the `audit_logs` cluster parameter is enabled, you must also enable Amazon DocumentDB to export logs to Amazon CloudWatch. If you omit either of these steps, audit logs will not be sent to CloudWatch.

Do not get confused with "Amazon DocumentDB records Data Definition Language (DDL), authentication, authorization, and user management events to Amazon CloudWatch Logs." The above event start to get recorded when you enable auditing via enabling `audit_log` parameter.

upvoted 5 times

🗨️ 👤 **kush_sumit** 2 years, 10 months ago

Answer A:

<https://docs.aws.amazon.com/documentdb/latest/developerguide/event-auditing.html>

Step 2. Enable Amazon CloudWatch Logs Export

When the value of the `audit_logs` cluster parameter is enabled, you must also enable Amazon DocumentDB to export logs to Amazon CloudWatch. If you omit either of these steps, audit logs will not be sent to CloudWatch.

When creating a cluster, performing a point-in-time-restore, or restoring a snapshot, you can enable CloudWatch Logs by following these steps.

Do not get confused with "Amazon DocumentDB records Data Definition Language (DDL), authentication, authorization, and user management events to Amazon CloudWatch Logs." The above event start to get recorded when you enable auditing via enabling `audit_log` parameter.

upvoted 1 times

A company is looking to move an on-premises IBM Db2 database running AIX on an IBM POWER7 server. Due to escalating support and maintenance costs, the company is exploring the option of moving the workload to an Amazon Aurora PostgreSQL DB cluster. What is the quickest way for the company to gather data on the migration compatibility?

- A. Perform a logical dump from the Db2 database and restore it to an Aurora DB cluster. Identify the gaps and compatibility of the objects migrated by comparing row counts from source and target tables.
- B. Run AWS DMS from the Db2 database to an Aurora DB cluster. Identify the gaps and compatibility of the objects migrated by comparing the row counts from source and target tables.
- C. Run native PostgreSQL logical replication from the Db2 database to an Aurora DB cluster to evaluate the migration compatibility.
- D. Run the AWS Schema Conversion Tool (AWS SCT) from the Db2 database to an Aurora DB cluster. Create a migration assessment report to evaluate the migration compatibility.

Suggested Answer: D

Reference:

<https://docs.aws.amazon.com/SchemaConversionTool/latest/userguide/Schema-Conversion-Tool.pdf>

Community vote distribution

D (100%)

🗳️ 👤 **LMax** Highly Voted 3 years, 6 months ago

Must be D

upvoted 9 times

🗳️ 👤 **adelcold** Most Recent 1 year, 11 months ago

Selected Answer: D

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/migrate-from-ibm-db2-on-amazon-ec2-to-aurora-postgresql-compatible-using-aws-dms-and-aws-sct.html>

upvoted 2 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: D

run AWS SCT from the Db2 database to an Aurora DB cluster.

-> Create a migration assessment report to evaluate the migration compatibility.

upvoted 2 times

🗳️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: D

SCT can do the required checks

upvoted 4 times

🗳️ 👤 **myutran** 3 years, 7 months ago

Ans: D

upvoted 2 times

🗳️ 👤 **JobinAkaJoe** 3 years, 7 months ago

D - right answer

upvoted 1 times

🗳️ 👤 **firbhat** 3 years, 7 months ago

ANS: D

- Converts DB/DW schema from source to target (including procedures / views / secondary indexes / FK and constraints)
- Mainly for heterogeneous DB migrations and DW migrations

upvoted 3 times

🗳️ 👤 **jnassp1** 3 years, 8 months ago

D - SCT

upvoted 2 times

An ecommerce company is using Amazon DynamoDB as the backend for its order-processing application. The steady increase in the number of orders is resulting in increased DynamoDB costs. Order verification and reporting perform many repeated GetItem functions that pull similar datasets, and this read activity is contributing to the increased costs. The company wants to control these costs without significant development efforts.


How should a Database Specialist address these requirements?

- A. Use AWS DMS to migrate data from DynamoDB to Amazon DocumentDB
- B. Use Amazon DynamoDB Streams and Amazon Kinesis Data Firehose to push the data into Amazon Redshift
- C. Use an Amazon ElastiCache for Redis in front of DynamoDB to boost read performance
- D. Use DynamoDB Accelerator to offload the reads

Suggested Answer: B

Community vote distribution

D (100%)

 **learnaws** Highly Voted 3 years, 8 months ago

Keywords here are "without significant development efforts". DAX is the answer. D
upvoted 17 times

 **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: D

D. Use DynamoDB Accelerator to offload the reads

https://docs.amazonaws.cn/en_us/amazondynamodb/latest/developerguide/DAX.html

upvoted 1 times

 **backbencher2022** 2 years, 2 months ago


Selected Answer: D

D is the correct answer. Per AWS document -

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/V2globaltables_HowItWorks.html

For read-heavy or bursty workloads, DAX provides increased throughput and potential operational cost savings by reducing the need to overprovision read capacity units. This is especially beneficial for applications that require repeated reads for individual keys.

upvoted 3 times

 **Arun32** 2 years, 6 months ago

Option D :

https://docs.amazonaws.cn/en_us/amazondynamodb/latest/developerguide/DAX.html

"Applications that are read-intensive, but are also cost-sensitive. With DynamoDB, you provision the number of reads per second that your application requires. If read activity increases, you can increase your tables' provisioned read throughput (at an additional cost). Or, you can offload the activity from your application to a DAX cluster, and reduce the number of read capacity units that you need to purchase otherwise."

upvoted 2 times

 **novice_expert** 3 years ago

Selected Answer: D

large number of repeated GetItem methods => DAX

upvoted 3 times

 **user0001** 3 years, 2 months ago

D is the answer since they don't want developments

upvoted 1 times

 **jove** 3 years, 5 months ago

Option D. Service cost is not a concern but development cost is.

upvoted 2 times

 **GMartinelli** 3 years, 6 months ago

Selected Answer: D

The problem is various repeated operations, a in memory DB like DAX is the best aproach.

upvoted 3 times

🗳️ 👤 **LMax** 3 years, 6 months ago

D is the easiest option

upvoted 2 times

🗳️ 👤 **myutran** 3 years, 6 months ago

Ans: D

upvoted 1 times

🗳️ 👤 **JobinAkaJoe** 3 years, 6 months ago

We need caching solution to offload reads and reduce the cost.

DAX is the best caching solution for DynamoDB API calls

upvoted 1 times

🗳️ 👤 **Ashoks** 3 years, 6 months ago

yes it is D

upvoted 1 times

🗳️ 👤 **AWSCert2020** 3 years, 6 months ago

D here, but DAX is not cost effective way

upvoted 2 times

🗳️ 👤 **[Removed]** 3 years, 7 months ago

B was never in play for me. Redshift is not cheap and that's not what they're asking.

I'm surprised nobody picked C, textbook case for Caching

upvoted 1 times

🗳️ 👤 **cloud4gr8** 3 years, 7 months ago

C is not easy to set up....it needs significant development efforts.

So, its not the ideal one choose.

D could be the best option.

upvoted 1 times

🗳️ 👤 **Dip11** 3 years, 6 months ago

C and D both caching solution. D is purpose built for DynamoDB, where in C you have to do deal with staleness manually. So D is correct answer.

upvoted 1 times

🗳️ 👤 **szmulder** 3 years, 7 months ago

Q, The company wants to control these costs

Answer B. I this to using Redshift and kinesis will increase the cost a lot compare to just using DAX.

So the answer is D

upvoted 2 times

🗳️ 👤 **BillyMadison** 3 years, 7 months ago

I think D as well. Dax helps with read intensiveness and cost effectiveness "without significant development efforts".

https://docs.amazonaws.cn/en_us/amazondynamodb/latest/developerguide/DAX.html

"Applications that are read-intensive, but are also cost-sensitive. With DynamoDB, you provision the number of reads per second that your application requires. If read activity increases, you can increase your tables' provisioned read throughput (at an additional cost). Or, you can offload the activity from your application to a DAX cluster, and reduce the number of read capacity units that you need to purchase otherwise."

upvoted 2 times

🗳️ 👤 **BillyC** 3 years, 8 months ago

D here

upvoted 4 times

An IT consulting company wants to reduce costs when operating its development environment databases. The company's workflow creates multiple Amazon

Aurora MySQL DB clusters for each development group. The Aurora DB clusters are only used for 8 hours a day. The DB clusters can then be deleted at the end of the development cycle, which lasts 2 weeks.

Which of the following provides the MOST cost-effective solution?

- A. Use AWS CloudFormation templates. Deploy a stack with the DB cluster for each development group. Delete the stack at the end of the development cycle.
- B. Use the Aurora DB cloning feature. Deploy a single development and test Aurora DB instance, and create clone instances for the development groups. Delete the clones at the end of the development cycle.
- C. Use Aurora Replicas. From the master automatic pause compute capacity option, create replicas for each development group, and promote each replica to master. Delete the replicas at the end of the development cycle.
- D. Use Aurora Serverless. Restore current Aurora snapshot and deploy to a serverless cluster for each development group. Enable the option to pause the compute capacity on the cluster and set an appropriate timeout.

Suggested Answer: D

Community vote distribution

D (61%)

B (39%)

 **Balki** Highly Voted 2 years, 11 months ago

Selected Answer: D

It is a close call between B and D. However, not everyone talks about the actual API call that can be used for Aurora Serverless, which is called as AutoPause. More details in the link.

https://docs.aws.amazon.com/AmazonRDS/latest/APIReference/API_ScalingConfiguration.html

So it should be D

upvoted 10 times

 **JobinAkaJoe** Highly Voted 3 years, 7 months ago

A - meets the basic requirement. But its not cost effective.

B - Good option. Saves cost on the storage layer with copy-on-write feature

C - Meets the requirement, bit not cost effective.

D - Most cost effective.


<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless.how-it-works.html#aurora-serverless.how-it-works.pause-resume>

upvoted 6 times

 **Dip11** 3 years, 6 months ago

Spot on

upvoted 1 times

 **MultiAZ** Most Recent 1 year, 4 months ago

Selected Answer: D

I go for D.


B looks promising. However, the requirement does not state the data in the clusters is from the same source. Also, keeping the DBs up only 8 hours a day is great for Serverless with auto-pause.

upvoted 1 times

 **aws2023a** 1 year, 9 months ago

A is also viable solution.

upvoted 1 times

 **aviathor** 1 year, 12 months ago

Selected Answer: D

Definitely D since that is the only option that also allows to save on compute by stopping unneeded resources outside of the 8 hours/5 days a week they are needed

upvoted 1 times

🗨️ **manig** 2 years ago

- 'Most Cost effective' is the Key
- Using clone db you can save on the storage cost but still you have to pay for Compute resources which are going to be charged more than 8Hrs.
- Aurora server less reduces the instance cost which will be the most cost saving

upvoted 2 times

🗨️ **RBSK** 2 years, 7 months ago

Selected Answer: B

Most-Cost effective solution is the key here. REfer to <https://aws.plainenglish.io/aurora-database-clones-what-they-are-and-when-to-use-them-b82be9d60309>

Pricing

Now let's talk about the pricing for creating and using Amazon Aurora database clones. We get billed per hour for the database instances provisioned as part of the clone database. The pricing is based on the DB instance class that we select for the clone. We are also charged for any addition storage that we use as we make edits to the clone database.

Option D is full copy of the DB and definitely will cost more than Option B.

upvoted 1 times

🗨️ **im_not_robot** 2 years, 3 months ago

storage cost is much lower than server cost, so option D is cheaper than B

upvoted 2 times

🗨️ **Chirantan** 2 years, 10 months ago

Answer should be D

upvoted 4 times

🗨️ **ryuhei** 2 years, 11 months ago

Selected Answer: D

Answer:D

upvoted 3 times

🗨️ **Radhaghosh** 3 years ago

Selected Answer: B

Option B is correct.

Question talks about destroying the cluster after 2 weeks, With aurora you will not pay for the compute capacity but you will pay for the storage capacity. If Option D had the delete option, I could have happily choose Option D. The advantage of Option B, you are not paying additional Storage as clone use the same storage + delta changes.

upvoted 1 times

🗨️ **[Removed]** 3 years ago

Selected Answer: D

It is necessary to support operation for 8 hours.

upvoted 2 times

🗨️ **novice_expert** 3 years ago

Selected Answer: B

single dev & test instance

-> Aurora clone

-> delete clones at end of dev cycle

upvoted 2 times

🗨️ **randss** 3 years, 1 month ago

Even if u pause serverless cluster storage, doesnt address the destroyed after 2 weeks requirement. I go with B.

upvoted 2 times

🗨️ **tugboat** 3 years, 3 months ago

Selected Answer: B

Amazon Aurora now allows you to create clones between Aurora Serverless v1 and provisioned Aurora DB clusters to enable quick sharing of data.



upvoted 4 times

🗨️ **awsmonster** 3 years, 3 months ago

I vote for B:

There isn't an "option to pause the compute capacity on the cluster and set an appropriate timeout." in the RDS console for Serverless.

upvoted 1 times



  **Shunpin** 3 years, 5 months ago

Selected Answer: B

I prefer B.

Aurora Serverless is not compatible to all Aurora provisioned engine version. However, you can do clone with most engine version. Meanwhile, I also consider the performance while restoring snapshot to Aurora Serverless.

upvoted 3 times

  **jove** 3 years, 5 months ago

It is D

upvoted 1 times

A company has multiple applications serving data from a secure on-premises database. The company is migrating all applications and databases to the AWS

Cloud. The IT Risk and Compliance department requires that auditing be enabled on all secure databases to capture all log ins, log outs, failed logins, permission changes, and database schema changes. A Database Specialist has recommended Amazon Aurora MySQL as the migration target, and leveraging the Advanced

Auditing feature in Aurora.

Which events need to be specified in the Advanced Auditing configuration to satisfy the minimum auditing requirements? (Choose three.)

- A. CONNECT
- B. QUERY_DCL
- C. QUERY_DDL
- D. QUERY_DML
- E. TABLE
- F. QUERY

Suggested Answer: ACE

Reference:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Auditing.html>

Enabling advanced auditing

Use the parameters described in this section to enable and configure Advanced Auditing for your DB cluster.

Use the `server_audit_logging` parameter to enable or disable Advanced Auditing, and the `server_audit_events` parameter to specify what events to log.

Use the `server_audit_excl_users` and `server_audit_incl_users` parameters to specify who gets audited. If `server_audit_excl_users` and `server_audit_incl_users` are empty (the default), all users are audited. If you add users to `server_audit_incl_users` and leave `server_audit_excl_users` empty, then only those users are audited.

Community vote distribution



alwaysAstudent Highly Voted 3 years, 6 months ago

Answer A,B,C

Connect - logins / DCL - authorizations (grant,revoke), DDL - schema updates

upvoted 15 times

johnconnor 3 years, 6 months ago

you are right

upvoted 1 times

user0001 3 years, 2 months ago

B is wrong , they are not asking for grant and revoke

upvoted 1 times

jthuma 2 years, 5 months ago

I think they are...Grant Revoke are permission statements.

upvoted 2 times

🗄️ 👤 **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: ABC

A. CONNECT

B. QUERY_DCL

C. QUERY_DDL

upvoted 2 times

🗄️ 👤 **IhorK** 1 year, 9 months ago

Selected Answer: ABC

- CONNECT – Logs both successful and failed connections and also disconnections. This event includes user information.

- QUERY_DCL – Similar to the QUERY event, but returns only data control language (DCL) queries (GRANT, REVOKE, and so on).

- QUERY_DDL – Similar to the QUERY event, but returns only data definition language (DDL) queries (CREATE, ALTER, and so on).

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Auditing.html>

upvoted 3 times

🗄️ 👤 **ninjalight25** 2 years, 2 months ago

Selected Answer: ABD

The events that need to be specified in the Advanced Auditing configuration to satisfy the minimum auditing requirements are:

A. CONNECT - to capture all log ins and log outs.

B. QUERY_DDL - to capture permission changes and database schema changes.

D. QUERY_DML - to capture data manipulation operations such as INSERT, UPDATE and DELETE.

upvoted 1 times

🗄️ 👤 **Isio05** 1 year, 9 months ago

They're not asking about changes in data, thus DML is not needed. Instead DCL is required to track permission changes.

upvoted 1 times

🗄️ 👤 **novice_expert** 3 years ago

Selected Answer: ABC

A. CONNECT (for log ins and log outs, unsuccessful logins)

B. QUERY_DCL (for grants - authorization modifications)

C. QUERY_DDL (for schema changes)

x D. QUERY_DML

x E. TABLE

x F. QUERY

upvoted 4 times

🗄️ 👤 **novice_expert** 3 years ago

Selected Answer: ABC

QUERY_DCL for authorization changes

upvoted 1 times

🗄️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: ABC

ABC - pretty straightforward

upvoted 3 times

🗄️ 👤 **user0001** 3 years, 2 months ago

ACD

since they need to capture data modification

upvoted 1 times

🗄️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: ABC

B is necessary too

upvoted 2 times

🗄️ 👤 **jove** 3 years, 5 months ago

A,B,C are correct

upvoted 1 times

A gaming company has recently acquired a successful iOS game, which is particularly popular during the holiday season. The company has decided to add a leaderboard to the game that uses Amazon DynamoDB. The application load is expected to ramp up over the holiday season. Which solution will meet these requirements at the lowest cost?

- A. DynamoDB Streams
- B. DynamoDB with DynamoDB Accelerator
- C. DynamoDB with on-demand capacity mode
- D. DynamoDB with provisioned capacity mode with Auto Scaling

Suggested Answer: C

Reference:

https://aws.amazon.com/blogs/database/running-spiky-workloads-and-optimizing-costs-by-more-than-90-using-amazon-dynamodb-on-demand-capacity-mode/?nc1=b_rp

AWS Database Blog

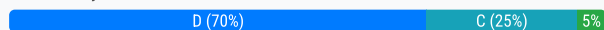
Running spiky workloads and optimizing costs by more than 90% using Amazon DynamoDB on-demand capacity mode

by Keisuke Utsumi | on 04 MAR 2020 | in [Amazon DynamoDB](#) | [Permalink](#) | [Comments](#) | [Share](#)

This is a guest post by Keisuke Utsumi, a Software Engineer with TVer Technologies Inc. In their own words, "TVer Technologies Inc. offers interactive entertainment services to users using a synchronized website with a TV broadcast."

TVer Technologies Inc. provides website and app-based interactive content for TV viewers in Japan. Many of our applications use [Amazon DynamoDB](#) as a database to store registered user information and logging the history of users' voting activity in live voting events during TV broadcasts. Our applications are often used by daily morning shows and seasonal pop music shows. In this blog post, we review how we have been able to optimize costs and performance for a system used in a TV live-voting event by using DynamoDB's on-demand read/write capacity mode.

Community vote distribution



Dantas Highly Voted 2 years, 12 months ago

Selected Answer: D

provisioned + auto scaling

1. It's not a new game, so the application's load (and its increase during Xmas) is already known.

2. the question asks "at the lowest possible cost":

upvoted 8 times

awsmonster Highly Voted 3 years, 4 months ago

I vote for D.

Reason, the game is bought from another company. So it has been running for some time and the load is known. The company is able to even identify the surge during Christmas period. Provisioned capacity will be more economical as compared to on demand.

upvoted 5 times

Germaneli Most Recent 1 year, 8 months ago

Selected Answer: C

On-Demand Capacity is the most suitable choice here.

The burst is expected only during holiday season, no other long-running changes are mentioned that would justify Provisioned Capacity. Also, a baseline for provisioned mode is not mentioned in the question, so there would be a reason for it and we can assume that provisioned mode has not been reasonable so far given the absence of that baseline.

Hence, On-Demand is the preference for this use case. Once the holidays are over (~weeks), everything is expected to go back to normal.

upvoted 1 times

Pranava_GCP 1 year, 8 months ago

Selected Answer: D

D. DynamoDB with provisioned capacity mode with Auto Scaling

should not be C. on-demand, "On-demand is ideal for bursty, new, or unpredictable workloads whose traffic can spike in seconds or minutes, and when underprovisioned capacity would impact the user experience."

<https://aws.amazon.com/blogs/database/amazon-dynamodb-auto-scaling-performance-and-cost-optimization-at-any-scale/>
upvoted 2 times

🗳️ 👤 **chen0305_099** 1 year, 8 months ago

Selected Answer: C

I THINK C

upvoted 1 times

🗳️ 👤 **manig** 1 year, 11 months ago

Answer: C

Auto-Scaling you are basically paying for throughput 24/7. Whereas for On-Demand Scaling you pay per request. This means for applications still in development or low traffic applications, it might be more economical to use On-Demand Scaling and not worry about provisioning throughput.

upvoted 2 times

🗳️ 👤 **redman50** 2 years, 2 months ago

Option D, DynamoDB with provisioned capacity mode with Auto Scaling, requires the company to provision a certain amount of read and write capacity ahead of time. While Auto Scaling can help with managing the scaling of the table based on demand, it comes with the additional cost of paying for the provisioned capacity even when it is not in use.

Therefore, Option C, DynamoDB with on-demand capacity mode, is the most cost-effective option that can meet the requirements of adding a leaderboard to the game and handling the expected increase in application load over the holiday season.

upvoted 1 times

🗳️ 👤 **backbencher2022** 2 years, 2 months ago

Selected Answer: D

Option D is a clear winner based on this blog - <https://aws.amazon.com/blogs/database/amazon-dynamodb-auto-scaling-performance-and-cost-optimization-at-any-scale/>

upvoted 2 times

🗳️ 👤 **ninjalight25** 2 years, 2 months ago

Selected Answer: C

The solution that will meet these requirements at the lowest cost is DynamoDB with on-demand capacity mode. This is because with on-demand capacity mode, DynamoDB automatically scales capacity up and down based on the actual request volume and traffic patterns.

upvoted 1 times

🗳️ 👤 **Jokera** 2 years, 8 months ago

Selected Answer: D

I vote for D

upvoted 2 times

🗳️ 👤 **RyanFeng** 2 years, 8 months ago

C might be better

<https://dynobase.dev/dynamodb-on-demand-vs-provisioned-scaling/>

upvoted 1 times

🗳️ 👤 **sachin** 2 years, 11 months ago

Since this is a leader board, DAX is out of question which mostly help in read intensive workloads.

Now when we compare on-Demand vs provisioned with auto scaling .

Since the application is already a popular application and likely to increase during christmas that mean business is aware about the application usage pattern.

<https://aws.amazon.com/blogs/database/amazon-dynamodb-auto-scaling-performance-and-cost-optimization-at-any-scale/>

Provisioned Auto Scaling can provide lower cost also in shown in the above case study.

upvoted 2 times



🗳️ 👤 **rlnd2000** 2 years, 11 months ago

Selected Answer: B

I will go with B.

The app is ok and the new feature is the leaderboard, leaderboard needs low latency so DAX is the answer for me.

upvoted 1 times

  **rind2000** 2 years, 11 months ago

One more thing DAX saves costs because it reduces read load in DynamoDB,

upvoted 1 times

  **novice_expert** 3 years ago

Selected Answer: C

x A. DynamoDB Streams

x B. DynamoDB with DynamoDB Accelerator (DAX is for fast static reads, but here data will change quickly)

C. DynamoDB with on-demand capacity mode

x D. DynamoDB with provisioned capacity mode with Auto Scaling (may be good option if game in use always)

upvoted 1 times

  **user0001** 3 years, 2 months ago

why not B ? it is about leaderboard refresh

upvoted 2 times

  **Germaneli** 1 year, 7 months ago

No, it's about a surge in writing - not reading - users. DAX no help here.

upvoted 1 times

  **jove** 3 years, 5 months ago

Auto Scaling is cheaper.. If you're not expecting sudden spikes. I vote for D

upvoted 2 times

  **grekh001** 3 years, 5 months ago

I'm not sure about C?

"On-demand is ideal for bursty, new, or unpredictable workloads whose traffic can spike in seconds or minutes"

vs.

"DynamoDB released auto scaling to make it easier for you to manage capacity efficiently, and auto scaling continues to help DynamoDB users lower the cost of workloads that have a predictable traffic pattern."

<https://aws.amazon.com/blogs/database/amazon-dynamodb-auto-scaling-performance-and-cost-optimization-at-any-scale/>

To save costs and based on the link I'd say that the answer is D

upvoted 3 times

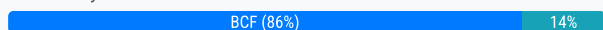
A company's Security department established new requirements that state internal users must connect to an existing Amazon RDS for SQL Server DB instance using their corporate Active Directory (AD) credentials. A Database Specialist must make the modifications needed to fulfill this requirement.

Which combination of actions should the Database Specialist take? (Choose three.)

- A. Disable Transparent Data Encryption (TDE) on the RDS SQL Server DB instance.
- B. Modify the RDS SQL Server DB instance to use the directory for Windows authentication. Create appropriate new logins.
- C. Use the AWS Management Console to create an AWS Managed Microsoft AD. Create a trust relationship with the corporate AD.
- D. Stop the RDS SQL Server DB instance, modify it to use the directory for Windows authentication, and start it again. Create appropriate new logins.
- E. Use the AWS Management Console to create an AD Connector. Create a trust relationship with the corporate AD.
- F. Configure the AWS Managed Microsoft AD domain controller Security Group.

Suggested Answer: CDF

Community vote distribution



manan728 Highly Voted 3 years, 6 months ago

This question was asked in my exam. B,C and F seems the correct options.

upvoted 6 times

IhorK Most Recent 1 year, 9 months ago

C. ...to create an AWS Managed Microsoft AD. Create a trust relationship...

E. ... to create an AD Connector. Create a trust relationship...

Both option are correct.

We should use AWS Managed Microsoft AD when have > 5000 users.

We should use AD Connector or Simple AD when have < 5000 users.

The question is not complete, or one of these answers should be excluded.

https://docs.aws.amazon.com/directoryservice/latest/admin-guide/ad_connector_best_practices.html

upvoted 2 times

adelcold 1 year, 11 months ago

Selected Answer: BCF

<https://aws.amazon.com/blogs/security/how-to-enable-windows-integrated-authentication-for-rds-for-sql-server-using-on-premises-active-directory/#:~:text=The%20setup%201%20Step%201%3A%20Set%20up%20RDS,between%20your%20VPC%20domain%20and%20your%20on-premises%20domain>

upvoted 1 times

ninjalight25 2 years, 2 months ago

Selected Answer: BCD

To fulfill the Security department's requirement for internal users to connect to the Amazon RDS for SQL Server DB instance using their corporate Active Directory credentials

upvoted 1 times

Mintwater 2 years, 1 month ago

D and B are conflicted, and B is the correct sequence -- modify the instance while the instance is up, let it connect to the AWS AD, then reboot the instance to make the connection valid. D -- Once you stopped the instance, you could not modify it.

upvoted 1 times

Mintwater 2 years, 1 month ago

BCF is the answer

upvoted 1 times

ryuhei 2 years, 11 months ago

Selected Answer: BCF

Answer:BCF



upvoted 1 times

  **novice_expert** 3 years ago

Selected Answer: BCF

- x A. Disable Transparent Data Encryption (unrelated)
- B. Modify the RDS SQL Server DB instance to use the directory for Windows authentication. Create appropriate new logins.(would need reboot)
- C. Use the AWS Management Console to create an AWS Managed Microsoft AD. Create a trust relationship with the corporate AD.
- x D. Stop the RDS SQL Server DB instance, modify it to use the directory for Windows authentication, and start it again. Create appropriate new logins.
(stop-start should be at end, not first)
- x E. Use the AWS Management Console to create an AD Connector. Create a trust relationship with the corporate AD.
- F. Configure the AWS Managed Microsoft AD domain controller Security Group.

upvoted 4 times

  **guru_ji** 3 years, 6 months ago

I got this Question in exam.

upvoted 1 times

  **myutran** 3 years, 6 months ago

Ans: BCF

upvoted 2 times

  **JobinAkaJoe** 3 years, 6 months ago

BCF

C&F is confirmed.



Choosing B over D because modifying the RDS to enable windows authentication must be done when the RDS is in available status though it will be rebooted for it to take effect.

upvoted 4 times

  **[Removed]** 3 years, 7 months ago

BCF is the answer - https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_SQLServerWinAuth.html

upvoted 3 times

  **Ashoks** 3 years, 7 months ago

BCF.

No restart required. Connector is a proxy, no trust relationship can be established with it.

upvoted 4 times

  **Manmohan** 3 years, 7 months ago

Ans: BEF


upvoted 1 times

  **anon9002** 3 years, 6 months ago

E - Not right

https://docs.aws.amazon.com/directoryservice/latest/admin-guide/ad_connector_app_compatibility.html

upvoted 2 times



  **Jack86** 3 years, 7 months ago

Sorry i made a mistake.

According to https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.DBInstance.Modifying.html#USER_ModifyInstance.Settings , modifying domain or directory id parameter in AWS RDS SqlServer requires a brief outage .

So correct answers are : C , D , F .

upvoted 4 times

  **Jack86** 3 years, 7 months ago

I would also vote for BCF.



Nobody doubt that C and F are correct.

Regarding my choice B (B or D ?).

We need to configure an existing DB (stated in the question).


According to <https://aws.amazon.com/blogs/database/joining-your-amazon-rds-instances-across-accounts-to-a-single-shared-domain/> , there is no need to stop the RDS to join it to an AD (step 3)

upvoted 4 times

  **AWSCert2020** 3 years, 7 months ago


CDF Here https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_SQLServerWinAuth.html

upvoted 3 times

  **Ebi** 3 years, 8 months ago

Answer is BDE

upvoted 1 times

  **Ebi** 3 years, 7 months ago

Meant BCF

D is not correct, you can't modify stopped RDS instance

upvoted 8 times

  **saki0915** 3 years, 7 months ago

Isn't B not restarting RDS?

upvoted 1 times

  **anon9002** 3 years, 6 months ago

E - Not right

https://docs.aws.amazon.com/directoryservice/latest/admin-guide/ad_connector_app_compatibility.html

upvoted 1 times

  **Mintwater** 2 years, 1 month ago

Note

Amazon RDS is compatible with AWS Managed Microsoft AD only, and is not compatible with AD Connector. For more information, see the AWS Microsoft AD section in the AWS Directory Service FAQs page.

upvoted 2 times

  **Billhardy** 3 years, 8 months ago

I will go with CDF

<https://www.powerupcloud.com/integrate-aws-sql-server-rds-with-multiple-ad/>

<https://www.sqlshack.com/advanced-windows-authentication-configurations-in-aws-rds-sql-server/>

<https://aws.amazon.com/blogs/security/how-to-enable-windows-integrated-authentication-for-rds-for-sql-server-using-on-premises-active-directory/>

<https://aws.amazon.com/blogs/aws/amazon-rds-for-sql-server-support-for-windows-authentication/>

upvoted 3 times

A Database Specialist is performing a proof of concept with Amazon Aurora using a small instance to confirm a simple database behavior. When loading a large dataset and creating the index, the Database Specialist encounters the following error message from Aurora:

ERROR: cloud not write block 7507718 of temporary file: No space left on device

What is the cause of this error and what should the Database Specialist do to resolve this issue?

- A. The scaling of Aurora storage cannot catch up with the data loading. The Database Specialist needs to modify the workload to load the data slowly.
- B. The scaling of Aurora storage cannot catch up with the data loading. The Database Specialist needs to enable Aurora storage scaling.
- C. The local storage used to store temporary tables is full. The Database Specialist needs to scale up the instance.
- D. The local storage used to store temporary tables is full. The Database Specialist needs to enable local storage scaling.

Suggested Answer: C

Reference:

<https://serverfault.com/questions/109828/how-can-i-tune-postgres-to-avoid-this-error>

how can i tune postgres to avoid this error?

Asked 11 years, 8 months ago Active 3 years, 5 months ago Viewed 23k times

I'm getting the error

7

ERROR: could not write block 3478284 of temporary file: No space left on device



when running the following query:



```
INSERT INTO summary SELECT t1.a, t1.b, SUM(t1.p) AS p, COUNT(t1.*) AS c,
    t1.d, t1.r, DATE_TRUNC('month', t1.start) AS month, t2.t AS t, t2.h, t2.x
FROM raw1 t1, raw2 t2
WHERE t1.t2_id=t2.id AND (t2.t<>'a' OR t2.y) GROUP BY month, t, a, b, d, r, h, x
```

table t1 is very large, and table t2 is pretty large

Caused by: org.postgresql.util.PSQLException: ERROR: could not write block 3478284 of temporary file: No space left on device
at org.postgresql.core.v3.QueryExecutorImpl.receiveErrorResponse(QueryExecutorImpl

Community vote distribution

C (100%)

novice_expert Highly Voted 3 years ago

Selected Answer: C

The local storage used to store temporary tables is full.

And only way to increase local storage is to upgrade instance type.

The Database Specialist needs to scale up the instance.

upvoted 6 times

jitesh_k Most Recent 1 year, 5 months ago

Is there something like local storage scaling? Talking about option D.

upvoted 1 times

Hisayuki 1 year, 4 months ago

Local storage is not defined by software, so what you can do for scaling is just scale up and changing the instance which has larger instance store. See also the link below:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Performance.html#AuroraMySQL.Managing.TempStorage>

upvoted 1 times



Pranava_GCP 1 year, 8 months ago

Selected Answer: C

C. The local storage used to store temporary tables is full. The Database Specialist needs to scale up the instance.

<https://repost.aws/knowledge-center/postgresql-aurora-storage-issue>



upvoted 1 times

  **adelcold** 1 year, 11 months ago

Selected Answer: C

<https://stackoverflow.com/questions/52610114/aurora-postgresql-engine-no-space-left-on-device>

upvoted 2 times

  **awsmonster** 3 years, 4 months ago



Answer: C

upvoted 2 times

  **user0001** 3 years, 2 months ago

true , you cant change local storage size , the only way is the change the instance type

upvoted 2 times

  **jove** 3 years, 5 months ago

Selected Answer: C

Option C

upvoted 2 times

  **GMartinelli** 3 years, 5 months ago

Selected Answer: C

Option C

upvoted 2 times

  **leunamE** 3 years, 6 months ago

Option C.

upvoted 3 times

A financial company wants to store sensitive user data in an Amazon Aurora PostgreSQL DB cluster. The database will be accessed by multiple applications across the company. The company has mandated that all communications to the database be encrypted and the server identity must be validated. Any non-SSL-based connections should be disallowed access to the database.

Which solution addresses these requirements?

- A. Set the `rds.force_ssl=0` parameter in DB parameter groups. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with `sslmode=allow`.
- B. Set the `rds.force_ssl=1` parameter in DB parameter groups. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with `sslmode=disable`.
- C. Set the `rds.force_ssl=0` parameter in DB parameter groups. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with `sslmode=verify-ca`.
- D. Set the `rds.force_ssl=1` parameter in DB parameter groups. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with `sslmode=verify-full`.

Suggested Answer: D

Reference:

<https://forums.aws.amazon.com/message.jspa?messageID=734076>

Community vote distribution

D (100%)

 **BillyC** Highly Voted 3 years, 8 months ago

ANS D is Correct!

upvoted 10 times

 **novice_expert** Highly Voted 3 years ago

Selected Answer: D

- in DB parameter groups: `rds.force_ssl=1` (`0=>false`, `1=>true`)
- Download and use the Amazon RDS certificate bundle
- configure the PostgreSQL connection string with `sslmode=verify-full`.

<https://jdbc.postgresql.org/documentation/head/ssl-client.html>

If `sslmode=verify-ca`, the server is verified by checking the certificate chain up to the root certificate stored on the client.

If `sslmode=verify-full`, the server host name will be verified to make sure it matches the name stored in the server certificate.

upvoted 6 times

 **khchan123** 3 years ago

Yes answer D.

upvoted 1 times

 **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: D

Answer is D

`rds.force_ssl=1` to force ssl in RDS and `sslmode=verify-full` to encrypt the connection and validate server identity.

upvoted 2 times


 **tucobbad** 2 years, 5 months ago

Selected Answer: D

Ans is D

`rds.force_ssl=1` to force ssl in RDS and `sslmode=verify-full` to encrypt the connection and validate server identity.

upvoted 3 times

 **guru_ji** 3 years, 6 months ago

Correct Answer ==> D

upvoted 1 times

🗨️ 👤 **myutran** 3 years, 6 months ago

Ans: D

upvoted 2 times

🗨️ 👤 **JobinAkaJoe** 3 years, 7 months ago

D indeed is the right choice

upvoted 1 times

🗨️ 👤 **Ashoks** 3 years, 7 months ago

yes. it is D

upvoted 2 times

🗨️ 👤 **firbhat** 3 years, 8 months ago

ANS: D

PostgreSQL: sslrootcert=rds-cert.pem sslmode=[verify-ca | verify-full]

upvoted 3 times

A company is using 5 TB Amazon RDS DB instances and needs to maintain 5 years of monthly database backups for compliance purposes. A Database

Administrator must provide Auditors with data within 24 hours.

Which solution will meet these requirements and is the MOST operationally efficient?

- A. Create an AWS Lambda function to run on the first day of every month to take a manual RDS snapshot. Move the snapshot to the company's Amazon S3 bucket.
- B. Create an AWS Lambda function to run on the first day of every month to take a manual RDS snapshot.
- C. Create an RDS snapshot schedule from the AWS Management Console to take a snapshot every 30 days.
- D. Create an AWS Lambda function to run on the first day of every month to create an automated RDS snapshot.

Suggested Answer: B

Community vote distribution



Exia Highly Voted 3 years, 7 months ago

A.

Unlike automated backups, manual snapshots aren't subject to the backup retention period. Snapshots don't expire.

For very long-term backups of MariaDB, MySQL, and PostgreSQL data, we recommend exporting snapshot data to Amazon S3. If the major version of your DB engine is no longer supported, you can't restore to that version from a snapshot.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_CreateSnapshot.html

upvoted 14 times

helpaws 2 years, 9 months ago

So the keyword here is "move to S3". Move means copy to S3 and then delete on RDS. The question asks for "operational efficiency" and "24 hours" to report to auditor. Just leave it in RDS, it doesn't expire and can be easily share right away. B is the answer.

upvoted 3 times

db2luwdba Highly Voted 2 years, 10 months ago

There is difference between copy / share a snapshot and Export. Export to S3 option will

When you export a DB snapshot, Amazon RDS extracts data from the snapshot and stores it in an Amazon S3 bucket. The data is stored in an Apache Parquet format that is compressed and consistent.

Where you have option to copy manual snapshot as is to different region or different AWS account. So we can not basically move the manual snapshot to S3 directly.

B is correct.

upvoted 10 times

db2luwdba 2 years, 10 months ago

Plus

Unlike automated backups, manual snapshots aren't subject to the backup retention period. Snapshots don't expire.

For very long-term backups of MariaDB, MySQL, and PostgreSQL data, we recommend exporting snapshot data to Amazon S3. If the major version of your DB engine is no longer supported, you can't restore to that version from a snapshot.

Here the backup movement is only for compliance. there is no requirement to query that backup .(using parquet format query through athena or rds redshift spectrum)

upvoted 3 times

michalf84 Most Recent 7 months, 1 week ago

Selected Answer: B

B no need to move snapshots as a suggests

upvoted 1 times

🗨️ 👤 **Michalf84** 8 months, 3 weeks ago

Selected Answer: B

No need to move the file it stays in AWS
upvoted 1 times

🗨️ 👤 **Rahul2406** 1 year, 1 month ago

B is correct
upvoted 1 times

🗨️ 👤 **MultiAZ** 1 year, 4 months ago

Selected Answer: B

B is the most operationally efficient.
upvoted 1 times

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

<https://aws.amazon.com/blogs/database/amazon-rds-snapshot-restore-and-recovery-demystified/#:~:text=Amazon%20RDS%20snapshots%20are%20stored,for%20copy%20and%20restore%20operations.>

Manual RDS snapshots are stored in S3 anyways. So moving to S3 does not make sense.
upvoted 1 times

🗨️ 👤 **Sathish_dbs** 1 year, 7 months ago

Selected Answer: B

Keep the snapshot in the RDS itself, no need to waste the operational efficiency by moving to and restoring from S3 unnecessarily
upvoted 1 times

🗨️ 👤 **Milan9527** 1 year, 11 months ago

Selected Answer: B

Move to S3 is additional work without efficiency
upvoted 1 times

🗨️ 👤 **Aviathor** 1 year, 12 months ago

Selected Answer: A

Actually you do not even need to take a manual snapshot. Even automated snapshots can be exported to S3.

You can export all types of DB snapshots—including manual snapshots, automated system snapshots, and snapshots created by the AWS Backup service.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ExportSnapshot.html

upvoted 1 times

🗨️ 👤 **Ninjalight25** 2 years, 2 months ago

Selected Answer: C

Option C, creating an RDS snapshot schedule from the AWS Management Console to take a snapshot every 30 days, would be the most operationally efficient solution for this scenario.

upvoted 1 times

🗨️ 👤 **Nice_Guy** 2 years, 2 months ago

I think the maximum retention period of automated backups is 35 days.

upvoted 2 times

🗨️ 👤 **Lollyj** 2 years, 5 months ago

Selected Answer: B

This is confusing because don't allow manual snapshots end up in S3 anyway?

upvoted 1 times

🗨️ 👤 **Maze** 2 years, 6 months ago

A.

Manual snapshot has a limitation. (Each supported Region: 100)

This case, customer wants to keep backup 5 years..

So I think, it can't be possible to keep snapshots during 5 years (365*5)

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Limits.html#RDS_Limits.Limits

upvoted 1 times

🗨️ 👤 **Maze** 2 years, 5 months ago

my mistake... B is the more correct.
because, only need to backups for a month..
so only need to snapshot $12 \times 5 = 60$.
it does not reach the rds snapshot's limitation.
upvoted 2 times

🗨️ 👤 **MrAliMohsan** 2 years ago

Also since they asked for "Operationally Efficient" Option so I also think B is a better answer.
upvoted 1 times

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

this limit is adjustable, you can really have more than 100 snaps per region
upvoted 1 times

🗨️ 👤 **awsjjj** 2 years, 7 months ago

Selected Answer: A

A is the answer. aws recommends long term backups to be exported to s3
upvoted 2 times

🗨️ 👤 **awsjjj** 2 years, 7 months ago

Selected Answer: A

For very long-term backups of MariaDB, MySQL, and PostgreSQL data, we recommend exporting snapshot data to Amazon S3. If the major version of your DB engine is no longer supported, you can't restore to that version from a snapshot.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_CreateSnapshot.html

upvoted 2 times

🗨️ 👤 **xyj** 2 years, 8 months ago

Option B is correct. manual snapshot won't expire.
upvoted 1 times

🗨️ 👤 **ryuhei** 2 years, 11 months ago

Selected Answer: B

I don't think you can move snapshots to individual S3, so the answer is probably B.
upvoted 4 times

A company wants to automate the creation of secure test databases with random credentials to be stored safely for later use. The credentials should have sufficient information about each test database to initiate a connection and perform automated credential rotations. The credentials should not be logged or stored anywhere in an unencrypted form.

Which steps should a Database Specialist take to meet these requirements using an AWS CloudFormation template?

- A. Create the database with the MasterUserName and MasterUserPassword properties set to the default values. Then, create the secret with the user name and password set to the same default values. Add a Secret Target Attachment resource with the SecretId and TargetId properties set to the Amazon Resource Names (ARNs) of the secret and the database. Finally, update the secret's password value with a randomly generated string set by the GenerateSecretString property.
- B. Add a Mapping property from the database Amazon Resource Name (ARN) to the secret ARN. Then, create the secret with a chosen user name and a randomly generated password set by the GenerateSecretString property. Add the database with the MasterUserName and MasterUserPassword properties set to the user name of the secret.
- C. Add a resource of type AWS::SecretsManager::Secret and specify the GenerateSecretString property. Then, define the database user name in the SecureStringTemplate template. Create a resource for the database and reference the secret string for the MasterUserName and MasterUserPassword properties. Then, add a resource of type AWS::SecretsManagerSecretTargetAttachment with the SecretId and TargetId properties set to the Amazon Resource Names (ARNs) of the secret and the database.
- D. Create the secret with a chosen user name and a randomly generated password set by the GenerateSecretString property. Add an SecretTargetAttachment resource with the SecretId property set to the Amazon Resource Name (ARN) of the secret and the TargetId property set to a parameter value matching the desired database ARN. Then, create a database with the MasterUserName and MasterUserPassword properties set to the previously created values in the secret.

Suggested Answer: C

Reference:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-secretsmanager-secrettargetattachment.html>

AWS::SecretsManager::SecretTargetAttachment

RSS

Filter View

All

The AWS::SecretsManager::SecretTargetAttachment resource completes the final link between a Secrets Manager secret and the associated database. This is required because each has a dependency on the other. No matter which one you create first, the other doesn't exist yet. To resolve this, you must create the resources in the following order:

1. Define the secret without referencing the service or database. You can't reference the service or database because it doesn't exist yet. The secret must contain a user name and password.

Community vote distribution

C (100%)

 **BillyC** Highly Voted 3 years, 8 months ago


Ans C is correct

upvoted 11 times

 **BillyMadison** 3 years, 7 months ago

BillyC, any idea why the AWS Database specialty exam is so hard to find on this site / why there are 404 errors?

upvoted 1 times

 **jyrajan** 3 years, 6 months ago

New requirement, only visible if you have contributor access

upvoted 1 times

🗲️ 👤 **IhorK** Most Recent 1 year, 9 months ago

Selected Answer: C

"Add a resource of type AWS::SecretsManager::RotationSchedule" missing in answer C.

upvoted 1 times

🗲️ 👤 **IhorK** 1 year, 9 months ago

Selected Answer: C

<https://malsouli.medium.com/aws-secrets-manager-create-and-rotate-secrets-automatically-36719faa7e4f>

upvoted 1 times

🗲️ 👤 **novice_expert** 3 years ago

Selected Answer: C

incomplete or wrong info but answer needs SecretsManager which is in C only

Add a resource of type AWS::SecretsManager::Secret

-> specify the GenerateSecretString property

-> define the database user name in the SecureStringTemplate template.

-> Create a resource for the database

-> reference the secret string for the MasterUserName and MasterUserPassword properties.

-> add a resource of type AWS::SecretsManagerSecretTargetAttachment with the SecretId and TargetId properties set to the Amazon Resource Names (ARNs) of the secret and the database.

upvoted 3 times

🗲️ 👤 **RotterDam** 3 years, 2 months ago

(C) is correct

upvoted 1 times

🗲️ 👤 **Ashoks** 3 years, 6 months ago

yes, it is C

upvoted 4 times

🗲️ 👤 **Ebi** 3 years, 7 months ago

Answer is C

upvoted 3 times

A company is going to use an Amazon Aurora PostgreSQL DB cluster for an application backend. The DB cluster contains some tables with sensitive data. A

Database Specialist needs to control the access privileges at the table level.

How can the Database Specialist meet these requirements?

- A. Use AWS IAM database authentication and restrict access to the tables using an IAM policy.
- B. Configure the rules in a NACL to restrict outbound traffic from the Aurora DB cluster.
- C. Execute GRANT and REVOKE commands that restrict access to the tables containing sensitive data.
- D. Define access privileges to the tables containing sensitive data in the pg_hba.conf file.

Suggested Answer: C

Reference:

<https://aws.amazon.com/blogs/database/managing-postgresql-users-and-roles/>

Community vote distribution

C (100%)

🗳️ 👤 **Pranava_GCP** 1 year, 9 months ago

Selected Answer: C

C. Execute GRANT and REVOKE commands that restrict access to the tables containing sensitive data.

[https://aws.amazon.com/blogs/database/managing-postgresql-users-and-](https://aws.amazon.com/blogs/database/managing-postgresql-users-and-roles/#:~:text=GRANT%20SELECT%20ON%20TABLE%20mytable1%2C%20mytable2%20TO%20readonly%3B)

[roles/#:~:text=GRANT%20SELECT%20ON%20TABLE%20mytable1%2C%20mytable2%20TO%20readonly%3B](https://aws.amazon.com/blogs/database/managing-postgresql-users-and-roles/#:~:text=GRANT%20SELECT%20ON%20TABLE%20mytable1%2C%20mytable2%20TO%20readonly%3B)

upvoted 2 times

🗳️ 👤 **IhorK** 1 year, 9 months ago

Selected Answer: C

<https://aws.amazon.com/blogs/database/managing-postgresql-users-and-roles/>

upvoted 2 times

🗳️ 👤 **sk1974** 2 years, 2 months ago

I am leaning towards A since you do not need to issue a GRANT to revoke a permission. U use grant to grant access to a table . C is kind of confusing

upvoted 2 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: C

x A. Use AWS IAM database authentication and restrict access to the tables using an IAM policy. (this is for db access)

x B. Configure the rules in a NACL to restrict outbound traffic from the Aurora DB cluster. (This is for Network Access Control)

C. Execute GRANT and REVOKE commands that restrict access to the tables containing sensitive data.

x D. Define access privileges to the tables containing sensitive data in the pg_hba.conf file. (not allowed)

upvoted 4 times

🗳️ 👤 **Zdujgr567783ff** 1 year, 12 months ago

revoke does not restrict access

it revokes some existing grant

upvoted 1 times

🗳️ 👤 **TonyGe** 3 years, 6 months ago

C.

Table level means DCL

upvoted 3 times

🗳️ 👤 **AM** 3 years, 7 months ago

This is easy. I am a DBA. Table level access is GRANT, REVOKE for many database flavors including Oracle and Postgres. Answer is C.

upvoted 2 times

🗳️ 👤 **myutran** 3 years, 7 months ago



Ans: C

upvoted 2 times

  **GeeBeeEl** 3 years, 7 months ago



I see answers like A or C, but I am looking for collateral support and cannot find any in the options selected so far. I choose A, see <https://aws.amazon.com/blogs/database/managing-postgresql-users-and-roles/> You create roles and then attach policies to the roles. May you support your response with a link so that we can check to confirm your reasoning? Thanks

upvoted 1 times

  **ExtHo** 3 years, 6 months ago



in your provided link roles are db roles not IAM roles :) so C is correct

upvoted 2 times

  **Ashoks** 3 years, 7 months ago

It is C

upvoted 3 times

  **Ebi** 3 years, 7 months ago


Answer is C

upvoted 1 times

  **Billhardy** 3 years, 7 months ago

This should be C

upvoted 2 times

  **BillyC** 3 years, 8 months ago

A or C.. please comments

upvoted 1 times

  **Kitty0403** 3 years, 8 months ago

Answer is C. Table level access is managed by DCL.

upvoted 7 times

A Database Specialist is working with a company to launch a new website built on Amazon Aurora with several Aurora Replicas. This new website will replace an on-premises website connected to a legacy relational database. Due to stability issues in the legacy database, the company would like to test the resiliency of Aurora.

Which action can the Database Specialist take to test the resiliency of the Aurora DB cluster?

- A. Stop the DB cluster and analyze how the website responds
- B. Use Aurora fault injection to crash the master DB instance
- C. Remove the DB cluster endpoint to simulate a master DB instance failure
- D. Use Aurora Backtrack to crash the DB cluster

Suggested Answer: B

Reference:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.FaultInjectionQueries.html>

Community vote distribution

B (100%)

🗳️ **TonyGe** Highly Voted 3 years, 6 months ago

B.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.FaultInjectionQueries.html>

upvoted 7 times

🗳️ **jitesh_k** Most Recent 1 year, 5 months ago

Cannot be A because stopping the DB cluster means application does not have a Database at all.

Cannot be C because removing DB cluster endpoint will also mean the same thing as above.

D is just a distractor - backtrack cannot crash a cluster.

B makes sense - it crashes one instance and allows Aurora to recover.

upvoted 1 times

🗳️ **Pranava_GCP** 1 year, 9 months ago

Selected Answer: B

B. Use Aurora fault injection to crash the master DB instance

upvoted 2 times

🗳️ **novice_expert** 3 years ago

Selected Answer: B

B. Use Aurora fault injection to crash the master DB instance

upvoted 4 times

🗳️ **myutran** 3 years, 6 months ago

Ans: B

upvoted 3 times

🗳️ **Ashoks** 3 years, 7 months ago

It is B

upvoted 3 times

🗳️ **BillyMadison** 3 years, 7 months ago

B with the link that is provided in the suggested answer

"You can test the fault tolerance of your Amazon Aurora DB cluster by using fault injection queries. Fault injection queries are issued as SQL commands to an Amazon Aurora instance and they enable you to schedule a simulated occurrence of one of the following events:

A crash of a writer or reader DB instance

A failure of an Aurora Replica



A disk failure

Disk congestion

When a fault injection query specifies a crash, it forces a crash of the Aurora DB instance. The other fault injection queries result in simulations of

failure events, but don't cause the event to occur. When you submit a fault injection query, you also specify an amount of time for the failure event simulation to occur for."

upvoted 4 times

  **firbhat** 3 years, 7 months ago

Ans B:

Two ways to test/simulate fault tolerance • Manual failover



- Fault injection queries

upvoted 3 times

  **Billhardy** 3 years, 8 months ago

will go with B

upvoted 1 times

  **BillyC** 3 years, 8 months ago

B is correct

upvoted 4 times

A company just migrated to Amazon Aurora PostgreSQL from an on-premises Oracle database. After the migration, the company discovered there is a period of time every day around 3:00 PM where the response time of the application is noticeably slower. The company has narrowed down the cause of this issue to the database and not the application.

Which set of steps should the Database Specialist take to most efficiently find the problematic PostgreSQL query?

- A. Create an Amazon CloudWatch dashboard to show the number of connections, CPU usage, and disk space consumption. Watch these dashboards during the next slow period.
- B. Launch an Amazon EC2 instance, and install and configure an open-source PostgreSQL monitoring tool that will run reports based on the output error logs.
- C. Modify the logging database parameter to log all the queries related to locking in the database and then check the logs after the next slow period for this information.
- D. Enable Amazon RDS Performance Insights on the PostgreSQL database. Use the metrics to identify any queries that are related to spikes in the graph during the next slow period.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **Pranava_GCP** 1 year, 8 months ago

Selected Answer: D

D. Enable Amazon RDS Performance Insights

"Whether your database performance problem is due to database configuration or application design issues, you can quickly identify the bottleneck and see which SQL statements are contributing to it."

<https://aws.amazon.com/rds/performance-insights/>
upvoted 2 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: D

RDS Performance Insights -> identify any queries that are related to spikes in the graph during the next slow period.
upvoted 3 times

🗳️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: D

Performance Insights
upvoted 2 times

🗳️ 👤 **sj143** 3 years, 3 months ago

if the database is transitioned to Aurora postgresSQL - how will enabling Amazon RDS performance insight help?
upvoted 1 times

🗳️ 👤 **VPup** 3 years, 3 months ago

D -

<https://aws.amazon.com/about-aws/whats-new/2021/10/rds-performance-insights-more-regions/>

" Amazon RDS Performance Insights is a database performance tuning and monitoring feature of RDS and Aurora that helps you quickly assess the load on your database and determine when and where to take action."

upvoted 4 times

🗳️ 👤 **guru_ji** 3 years, 6 months ago



Answer: D

upvoted 2 times

🗳️ 👤 **myutran** 3 years, 7 months ago

Ans: D

upvoted 1 times

  **Ashoks** 3 years, 7 months ago

D is ..

upvoted 2 times

  **BillyMadison** 3 years, 7 months ago

<https://aws.amazon.com/blogs/database/optimizing-and-tuning-queries-in-amazon-rds-postgresql-based-on-native-and-external-tools/>

You can monitor SQL queries that caused load, I/O waits, and the users and hosts through which the queries ran.

As you can see in the previous screenshot, the same SQL query has been reported as consuming the most CPUs.

Performance Insights supports counter metrics for RDS PostgreSQL. Counter metrics allow you to customize your Performance Insights dashboard to include up to 10 additional graphs from the available operating system and database metrics. It is helpful to identify and analyze performance issues by correlating load charts.

Performance counter metrics are native and non-native.

As you can see in the following screenshot, counter metrics are updated with tuples fetched, tuples returned, blocks latency, and blocks read.

upvoted 4 times

  **BillyMadison** 3 years, 7 months ago

Going with D for now because of <https://aws.amazon.com/blogs/database/optimizing-and-tuning-queries-in-amazon-rds-postgresql-based-on-native-and-external-tools/>

"AWS recently released a feature called Amazon RDS Performance Insights, which provides an easy-to-understand dashboard for detecting performance problems in terms of load."

"AWS recently released a feature called Amazon RDS Performance Insights, which provides an easy-to-understand dashboard for detecting performance problems in terms of load."

The other possible answers are B & C but those solutions include the possibility it is an application issue which the question says it is not.

upvoted 3 times

  **BillyC** 3 years, 7 months ago

Ans D is correct

upvoted 2 times

A company has a web-based survey application that uses Amazon DynamoDB. During peak usage, when survey responses are being collected, a Database Specialist sees the `ProvisionedThroughputExceededException` error. What can the Database Specialist do to resolve this error? (Choose two.)

- A. Change the table to use Amazon DynamoDB Streams
- B. Purchase DynamoDB reserved capacity in the affected Region
- C. Increase the write capacity units for the specific table
- D. Change the table capacity mode to on-demand
- E. Change the table type to throughput optimized

Suggested Answer: CE

Reference:

<https://forums.aws.amazon.com/thread.jspa?threadID=174315>

Community vote distribution

CD (100%)

pan24 Highly Voted 3 years, 8 months ago

Ans:CD

Dynamodb doesn't have a throughput optimized table type

upvoted 21 times

JobinAkaJoe Highly Voted 3 years, 7 months ago

- A. Change the table to use Amazon DynamoDB Streams --> Doesn't make any sense for the given problem.
 - B. Purchase DynamoDB reserved capacity in the affected Region --> reserved capacity is not going to give you anything more than you reserve
 - C. Increase the write capacity units for the specific table --> This could be an answer. But the exception is thrown when survey responses are collected which appears to be a read operation.
 - D. Change the table capacity mode to on-demand --> This is correct answer
- <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/switching.capacitymode.html>
- E. Change the table type to throughput optimized --> There is no such thing as throughput optimized in dynamoDB

I will go with C,D

upvoted 6 times

IhorK 1 year, 9 months ago

"survey responses are collected" - writing to DB, write capacity. C and D.

upvoted 2 times

Pranava_GCP Most Recent 1 year, 9 months ago

Selected Answer: CD

- C. Increase the write capacity units for the specific table
 - D. Change the table capacity mode to on-demand
- upvoted 1 times

lollyj 2 years, 5 months ago

Selected Answer: CD

My answer

upvoted 1 times

sriexam 2 years, 10 months ago

C and D - Any setting will be on Table level for DynamoDB

upvoted 1 times

minhntm 2 years, 11 months ago

Selected Answer: CD

Ans:CD

upvoted 1 times

🗨️ 👤 **dbaroger** 2 years, 12 months ago

C and D are not compatible. If you use on-demand mode, you don't change read/write capacity.

upvoted 1 times

🗨️ 👤 **praffuln** 3 years ago

ANS : C & D.

upvoted 1 times

🗨️ 👤 **novice_expert** 3 years ago

Selected Answer: CD

C. Increase the write capacity units for the specific table --> This could be an answer. As the exception is thrown when survey responses are collected which is write operation.

D. Change the table capacity mode to on-demand --> This is correct answer

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

upvoted 3 times

🗨️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: CD

CD

Whoever writes these questions should go for English grammar classes!! both options cannot go together which means a question like "What is database specialist's role in this?" makes little sense. What would have made sense is "What are the options a specialist can try to remediate this problem?"

upvoted 3 times

🗨️ 👤 **IhorK** 1 year, 9 months ago

Now imagine how difficult it is to understand everything that they want from you for a non-native speaker :)

upvoted 1 times

🗨️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: CD

C & D are the absolutely correct answers.

upvoted 1 times

🗨️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: CD

Correct answers

upvoted 1 times

🗨️ 👤 **kped21** 3 years, 3 months ago

C,D: Valid

upvoted 1 times

🗨️ 👤 **Dip11** 3 years, 6 months ago

C and D are the only correct options. Although both cannot go together.

Question does not clarify that both options should apply or either of it.

upvoted 1 times

🗨️ 👤 **myutran** 3 years, 6 months ago

Ans: CD

upvoted 3 times

🗨️ 👤 **MultiAZ** 3 years, 6 months ago

Should be C OR D, right... because either would work, but you cannot have both. On-demand does not allow you to set write capacity???

The only 2 answers that fit together are B and C, as both are about provisioned capacity. But B is very unlikely here, although may be a good practice overall.

upvoted 1 times

🗨️ 👤 **Ashoks** 3 years, 7 months ago

yes, C & D

upvoted 2 times

A company is running a two-tier ecommerce application in one AWS account. The web server is deployed using an Amazon RDS for MySQL Multi-AZ DB instance. A Developer mistakenly deleted the database in the production environment. The database has been restored, but this resulted in hours of downtime and lost revenue.

Which combination of changes in existing IAM policies should a Database Specialist make to prevent an error like this from happening in the future? (Choose three.)

- A. Grant least privilege to groups, users, and roles
- B. Allow all users to restore a database from a backup that will reduce the overall downtime to restore the database
- C. Enable multi-factor authentication for sensitive operations to access sensitive resources and API operations
- D. Use policy conditions to restrict access to selective IP addresses
- E. Use AccessList Controls policy type to restrict users for database instance deletion
- F. Enable AWS CloudTrail logging and Enhanced Monitoring

Suggested Answer: ACD

Community vote distribution

ACD (58%)

ACE (42%)

  **chicagomassageseeker** Highly Voted 3 years, 8 months ago

ACD are right

upvoted 14 times

  **bipin_john** 1 year, 1 month ago

Correct answer

upvoted 1 times

  **novice_expert** Highly Voted 3 years ago

Selected Answer: ACD

A. Least privileges

C. Multi factor

D. restrict access

upvoted 5 times

  **MultiAZ** Most Recent 1 year, 4 months ago

Selected Answer: ACE

I go fro ACE.

IP address is not a reliable decision tool.

upvoted 1 times

  **Pranava_GCP** 1 year, 9 months ago


Selected Answer: ACD

A. Grant least privilege to groups, users, and roles

C. Enable multi-factor authentication for sensitive operations to access sensitive resources and API operations

D. Use policy conditions to restrict access to selective IP addresses

upvoted 1 times

  **jiyakurani** 1 year, 11 months ago

ACD are correct options

upvoted 1 times

  **teo2157** 2 years, 2 months ago

Selected Answer: ACD

There's no AccessList Controls for RDS, so based on that goes for ACD



upvoted 2 times

  **ninjalight25** 2 years, 2 months ago

Selected Answer: ACD

the correct options are A, C, and D.

upvoted 1 times

  **lollyj** 2 years, 5 months ago

Selected Answer: ACE

Can someone explain why E isn't plausible? I didn't choose D because sometimes developers are need access to Prod environments and restricting their IPs doesn't mean they can't utilize another IP to do the same damage.

upvoted 3 times

  **Mintwater** 2 years, 1 month ago

E -- Access Control List -- not for RDS, IAM.

IAM have policies to manage the privileges, but no list

upvoted 1 times

  **khun** 2 years, 5 months ago

Selected Answer: ACE

ACE.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "rds:DeleteDBInstance"
      ],
      "Resource": [
        "arn:aws:rds::db:"
      ],
      "Effect": "Deny"
    }
  ]
}
```

upvoted 4 times

  **db2luwdba** 2 years, 10 months ago

Prevent a user from deleting a DB instance

The following permissions policy grants permissions to prevent a user from deleting a specific DB instance. For example, you might want to deny the ability to delete your production DB instances to any user that is not an administrator.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyDelete1",
      "Effect": "Deny",
      "Action": "rds:DeleteDBInstance",
      "Resource": "arn:aws:rds:us-west-2:123456789012:db:mysql-instance"
    }
  ]
}
```



ABE-- Wording of the question not very accurate. RDS Access polict can be done this way but there is nothing called as Access Control list poicy type in RDS

upvoted 3 times

  **db2luwdba** 2 years, 10 months ago

I mean ACE

upvoted 3 times

  **tugboat** 3 years, 3 months ago

Selected Answer: ACD

Appropriate and RDS supported options

upvoted 2 times

🗨️ 👤 **awsmonster** 3 years, 3 months ago

Agree with ACD

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/security_iam_id-based-policy-examples.html

upvoted 5 times

🗨️ 👤 **Mintwater** 2 years, 1 month ago

Policy best practices ---

Require multi-factor authentication (MFA)

Apply least-privilege permissions

Get started with AWS managed policies and move toward least-privilege permissions

Use conditions in IAM policies to further restrict access

upvoted 1 times

🗨️ 👤 **guru_ji** 3 years, 6 months ago

I got this Question in exam.

60% questions came in actual exam from this 145 set. Bunch of new Questions.

We can share study material for free, You can email me on

"awsdbguru at gmail"

upvoted 4 times

🗨️ 👤 **guru_ji** 3 years, 6 months ago

Correct Answer ==>> ACD

any idea how much Q we will get in real exam from Q available here?

anyone is preparing for this exam and want to do group study with us, comment with mail_id.

upvoted 1 times

🗨️ 👤 **manan728** 3 years, 7 months ago

ACD are correct choices. MFA is specified in the aws docs specifically for such use case

<https://aws.amazon.com/blogs/database/using-iam-multifactor-authentication-with-amazon-rds/>

upvoted 3 times

🗨️ 👤 **ricksun** 3 years, 7 months ago

I go for ACE

upvoted 4 times

🗨️ 👤 **ricksun** 3 years, 7 months ago

change to ACD since RDS not support access list

upvoted 1 times

🗨️ 👤 **myutran** 3 years, 7 months ago

Ans: ACD

upvoted 1 times

A company is building a new web platform where user requests trigger an AWS Lambda function that performs an insert into an Amazon Aurora MySQL DB cluster. Initial tests with less than 10 users on the new platform yielded successful execution and fast response times. However, upon more extensive tests with the actual target of 3,000 concurrent users, Lambda functions are unable to connect to the DB cluster and receive too many connections errors.

Which of the following will resolve this issue?

- A. Edit the my.cnf file for the DB cluster to increase max_connections
- B. Increase the instance size of the DB cluster
- C. Change the DB cluster to Multi-AZ
- D. Increase the number of Aurora Replicas

Suggested Answer: B

Community vote distribution

B (100%)

 **novice_expert** Highly Voted 3 years ago

Selected Answer: B

B. Increase the instance size of the DB cluster

we need more connections which depend on instance

Max_connection is a formula in RDS parameter group:

$\text{GREATEST}(\{\log(\text{DBInstanceClassMemory}/805306368)*45\}, \{\log(\text{DBInstanceClassMemory}/8187281408)*1000\})$

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Performance.html>

You can increase the maximum number of connections to your Aurora MySQL DB instance by scaling the instance up to a DB instance class with more memory, or by setting a larger value for the max_connections parameter in the DB parameter group for your instance, up to 16,000.

You must change a larger value for the max_connections parameter in the DB parameter group, not edit my.cnf, it is not physical server hosting MySQL.

upvoted 6 times

 **Pranava_GCP** Most Recent 1 year, 9 months ago


Selected Answer: B

B. Increase the instance size of the DB cluster

"You can increase the maximum number of connections to your Aurora MySQL DB instance by scaling the instance up to a DB instance class with more memory, or by setting a larger value for the max_connections parameter in the DB parameter group for your instance, up to 16,000."

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Performance.html>

upvoted 1 times


 **Shunpin** 3 years, 5 months ago

Selected Answer: B

Max_connection is a formula in RDS parameter group:

$\text{GREATEST}(\{\log(\text{DBInstanceClassMemory}/805306368)*45\}, \{\log(\text{DBInstanceClassMemory}/8187281408)*1000\})$

upvoted 3 times

 **guru_ji** 3 years, 7 months ago

I got this Question in exam.

Answer: B

upvoted 4 times

 **ChauPhan** 3 years, 7 months ago

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Performance.html>

You can increase the maximum number of connections to your Aurora MySQL DB instance by scaling the instance up to a DB instance class with

more memory, or by setting a larger value for the max_connections parameter in the DB parameter group for your instance, up to 16,000. You must change a larger value for the max_connections parameter in the DB parameter group, not edit my.cnf, it is not physical server hosting MySQL.

upvoted 3 times

🗨️ 👤 **Dip11** 3 years, 7 months ago

ANS B. RDS doesn't allow to change config files on host so A is ruled out.

upvoted 1 times

🗨️ 👤 **myutran** 3 years, 7 months ago

Ans: B

upvoted 2 times

🗨️ 👤 **JobinAkaJoe** 3 years, 8 months ago

B is the best answer

upvoted 2 times

🗨️ 👤 **Billhardy** 3 years, 8 months ago

B is correct

upvoted 2 times

🗨️ 👤 **BillyMadison** 3 years, 8 months ago

Agree with B because of <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Performance.html>

"You can increase the maximum number of connections to your Aurora MySQL DB instance by scaling the instance up to a DB instance class with more memory, or by setting a larger value for the max_connections parameter, up to 16,000."

Could be A as well, just haven't seen good enough documentation to choose it.

upvoted 4 times

🗨️ 👤 **ChauPhan** 3 years, 7 months ago

You can change a larger value for the max_connections parameter in the DB parameter group, not edit my.cnf, so A is incorrect

upvoted 3 times

🗨️ 👤 **BillyC** 3 years, 8 months ago

Ans B is correct

upvoted 4 times

A company is developing a multi-tier web application hosted on AWS using Amazon Aurora as the database. The application needs to be deployed to production and other non-production environments. A Database Specialist needs to specify different MasterUsername and MasterUserPassword properties in the AWS



CloudFormation templates used for automated deployment. The CloudFormation templates are version controlled in the company's code repository. The company also needs to meet compliance requirement by routinely rotating its database master password for production. What is most secure solution to store the master password?

- A. Store the master password in a parameter file in each environment. Reference the environment-specific parameter file in the CloudFormation template.
- B. Encrypt the master password using an AWS KMS key. Store the encrypted master password in the CloudFormation template.
- C. Use the secretsmanager dynamic reference to retrieve the master password stored in AWS Secrets Manager and enable automatic rotation.
- D. Use the ssm dynamic reference to retrieve the master password stored in the AWS Systems Manager Parameter Store and enable automatic rotation.

Suggested Answer: C

Community vote distribution

C (100%)

  **BillyMadison** Highly Voted 3 years, 7 months ago

Agree with C

"By using the secure string support in CloudFormation with dynamic references you can better maintain your infrastructure as code. You'll be able to avoid hard coding passwords into your templates and you can keep these runtime configuration parameters separated from your code. Moreover, when properly used, secure strings will help keep your development and production code as similar as possible, while continuing to make your infrastructure code suitable for continuous deployment pipelines."



<https://aws.amazon.com/blogs/mt/using-aws-systems-manager-parameter-store-secure-string-parameters-in-aws-cloudformation-templates/>

<https://aws.amazon.com/blogs/security/how-to-use-aws-secrets-manager-rotate-credentials-amazon-rds-database-types-oracle/>
upvoted 10 times

  **BillyC** Highly Voted 3 years, 8 months ago

C is correct

upvoted 7 times

  **ychaabane** Most Recent 1 year, 3 months ago

Selected Answer: C


agree with C

upvoted 1 times

  **Pranava_GCP** 1 year, 9 months ago

Selected Answer: C

C. Use the secretsmanager dynamic reference to retrieve the master password stored in AWS Secrets Manager and enable automatic rotation.
upvoted 2 times

  **ryuhei** 2 years, 11 months ago

Selected Answer: C

Answer:C

upvoted 1 times


  **novice_expert** 3 years ago

Selected Answer: C

C. Use the secretsmanager dynamic reference to retrieve the master password stored in AWS Secrets Manager and enable automatic rotation.

"By using the secure string support in CloudFormation with dynamic references you can better maintain your infrastructure as code. You'll be able to avoid hard coding passwords into your templates and you can keep these runtime configuration parameters separated from your code. Moreover, when properly used, secure strings will help keep your development and production code as similar as possible, while continuing to make your infrastructure code suitable for continuous deployment pipelines."

upvoted 2 times

  **selva1982** 3 years, 4 months ago

C is correct

upvoted 1 times

  **swarndeep** 3 years, 6 months ago



C is correct

upvoted 1 times

  **JobinAkaJoe** 3 years, 7 months ago



C is the right answe

upvoted 1 times

  **Ashoks** 3 years, 7 months ago

Ans - C



upvoted 2 times

  **Ebi** 3 years, 7 months ago

Answer id C

SSM does not supported automatic rotation

upvoted 6 times

  **BillyC** 3 years, 8 months ago

Sorry D

upvoted 1 times

A company is writing a new survey application to be used with a weekly televised game show. The application will be available for 2 hours each week. The company expects to receive over 500,000 entries every week, with each survey asking 2-3 multiple choice questions of each user. A Database Specialist needs to select a platform that is highly scalable for a large number of concurrent writes to handle the anticipated volume. Which AWS services should the Database Specialist consider? (Choose two.)

- A. Amazon DynamoDB
- B. Amazon Redshift
- C. Amazon Neptune
- D. Amazon Elasticsearch Service
- E. Amazon ElastiCache

Suggested Answer: AE

Community vote distribution

AE (100%)

  **BillyC**  3 years, 8 months ago

Ans A and E are correct

upvoted 8 times

  **guru_ji**  3 years, 6 months ago

I got this Question in exam.

upvoted 5 times

  **Pranava_GCP**  1 year, 9 months ago

Selected Answer: AE

A. Amazon DynamoDB

E. Amazon ElastiCache

upvoted 1 times

  **f__16** 2 years ago

My answer is AE too. But my question is can DynamoDB and Elasticache be used together? Are there many scenarios of using them both in the same application?

upvoted 2 times

  **novice_expert** 3 years ago

Selected Answer: AE

A. Amazon DynamoDB (for high responses 500K in 2 hours)

E. Amazon ElastiCache (for high static reads 500K in 2 hours for few questions + options)

upvoted 4 times

  **Shunpin** 3 years, 5 months ago

Selected Answer: AE

Advantages and disadvantages of write-through

<https://docs.aws.amazon.com/AmazonElastiCache/latest/mem-ug/Strategies.html#Strategies.WriteThrough>

upvoted 1 times

  **jove** 3 years, 5 months ago

Selected Answer: AE

<https://aws.amazon.com/products/databases/real-time-apps-elasticache-for-redis/>

upvoted 3 times

  **ChauPhan** 3 years, 6 months ago

it is asking about the use case of each DB

C. Amazon Neptune

It is not relevant, AWS Neptune is Graph database

D. Amazon Elasticsearch Service

it is not a use case, ES is using for log store/search/metric/config info/document list/ etc

E. Amazon ElastiCache

It is used to cache the data for faster query (read) performance, usually using before a database

A, B is my choice. Redshift is dataware house, but can be used as a data lake as high concurrent write.

upvoted 2 times

  **im_not_robot** 2 years, 3 months ago

redshift is not good with write-intensive use case

upvoted 1 times

  **RotterDam** 3 years, 2 months ago



definitely not Redshift

upvoted 4 times

  **jove** 3 years, 5 months ago

Redshift is not a good idea

upvoted 1 times

  **Aesthet** 3 years, 6 months ago

AE would make sense if not this part: "needs to select a platform that is highly scalable for a large number of concurrent writes".

I think AD is a better choice here.

upvoted 1 times

  **myutran** 3 years, 6 months ago

Ans: AE

upvoted 1 times

  **Exia** 3 years, 7 months ago

A, E for sure

<https://aws.amazon.com/products/databases/real-time-apps-elasticache-for-redis/>

upvoted 2 times

  **JobinAkaJoe** 3 years, 7 months ago

A. Amazon DynamoDB -- ideal for this requirement

B. Amazon Redshift -- Wrong choice

C. Amazon Neptune -- Not a requirement to have graph database

D. Amazon Elasticsearch Service -- Not requirement for ElasticSearch

E. Amazon ElastiCache -- The requirement is write intensive.. Not sure how Elasticache can help.



A,E seem to be the best choice.

upvoted 2 times

  **addixion** 3 years, 7 months ago

ElastiCache store the questions and multiple answers

upvoted 1 times

  **Ashoks** 3 years, 7 months ago

yes, A, E

upvoted 2 times

  **BillyMadison** 3 years, 7 months ago

AE

<https://aws.amazon.com/elasticache/>

Building real-time apps across versatile use cases like gaming, geospatial service, caching, session stores, or queuing, with advanced data structures, replication, and point-in-time snapshot support.

<https://aws.amazon.com/dynamodb/>

Build powerful web applications that automatically scale up and down. You don't need to maintain servers, and your applications have automated high availability.



upvoted 4 times

  **szmulder** 3 years, 7 months ago

My ans is A and D. Why not E, E is store data in memory and normally we use it as a buffer server but not the server to store the data.

D, Amazon Elasticsearch is Highly scalable.



upvoted 2 times



  **awsmonster** 3 years, 4 months ago

Amazon OpenSearch Service (successor to Amazon Elasticsearch Service)
makes it easy for you to perform interactive log analytics, real-time application monitoring, website search, and more.

Seems like an overkill for a 2-3 multiple question.

I vote for AE
upvoted 1 times

  **rootkim** 3 years, 7 months ago
My opinion is the same as yours.
upvoted 1 times

  **Ebi** 3 years, 7 months ago
Answer is AE
upvoted 1 times

A company has migrated a single MySQL database to Amazon Aurora. The production data is hosted in a DB cluster in VPC_PROD, and 12 testing environments are hosted in VPC_TEST using the same AWS account. Testing results in minimal changes to the test data. The Development team wants each environment refreshed nightly so each test database contains fresh production data every day.

Which migration approach will be the fastest and most cost-effective to implement?

- A. Run the master in Amazon Aurora MySQL. Create 12 clones in VPC_TEST, and script the clones to be deleted and re-created nightly.
- B. Run the master in Amazon Aurora MySQL. Take a nightly snapshot, and restore it into 12 databases in VPC_TEST using Aurora Serverless.
- C. Run the master in Amazon Aurora MySQL. Create 12 Aurora Replicas in VPC_TEST, and script the replicas to be deleted and re-created nightly.
- D. Run the master in Amazon Aurora MySQL using Aurora Serverless. Create 12 clones in VPC_TEST, and script the clones to be deleted and re-created nightly.

Suggested Answer: A

Community vote distribution

A (60%)

D (40%)

🗳️ 👤 **zanhsieh** Highly Voted 3 years, 8 months ago

A.
B dropped due to snapshot is slower (full disk dump) than clone (copy-on-write)
C dropped due to no write on Aurora Replicas
D dropped due to there's no option for cloning in the console.
upvoted 13 times

🗳️ 👤 **Huy** 3 years, 7 months ago

Cloning is not supported on Aurora Serverless nor Cross-Region.
<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Clone.html>
upvoted 5 times

🗳️ 👤 **ChauPhan** 3 years, 7 months ago

No, it is cross-account, not cross-region. The question mentions same account
You can create an Aurora provisioned clone from a provisioned Aurora DB cluster. You can create an Aurora Serverless v1 clone from an Aurora Serverless v1 DB cluster. But you can also create Aurora Serverless v1 clones from Aurora provisioned DB clusters, and you can create provisioned clones from Aurora Serverless v1 DB clusters.
"CROSS-ACCOUNT cloning currently doesn't support cloning Aurora Serverless v1 DB clusters"
upvoted 1 times

🗳️ 👤 **MultiAZ** Most Recent 1 year, 4 months ago

Selected Answer: A

Answer is A
The question does not mention variable workload, so I see no need for Serverless (D)
upvoted 1 times

🗳️ 👤 **IhorK** 1 year, 9 months ago

Selected Answer: A

Choose between A and D.
To have Aurora Serverless we need to convert RDS Aurora MySQL to Aurora Serverless MySQL.
We can do it 3 ways:
- Snapshot restore
- Logical backup and restore
- A new serverless reader (for Amazon Aurora PostgreSQL-Compatible Edition versions 13.6 and later). Add a serverless reader, force a failover. This promotes the reader instance to a writer instance.
<https://repost.aws/knowledge-center/aurora-migrate-provisioned-serverless>
We are asked for "the fastest" solution, so, answer A.
upvoted 2 times

🗳️ 👤 **Pankaj24hrs** 2 years ago

D

<https://aws.amazon.com/about-aws/whats-new/2021/06/amazon-aurora-serverless-v1-supports-fast-database-cloning/>

Aurora Serverless supports fast database cloning. You only pay for additional storage if you make data changes in the cloned DB cluster.

In question it mentioned "Testing results in minimal changes to the test data" so there will be a minimal cost for dev env databases. Most cost-effective.

upvoted 1 times

🗨️ 👤 **Mintwater** 2 years, 2 months ago

A.

it is possible to create an Aurora Cluster with a replica and then use it to create an Aurora Serverless cluster. Then use the Serverless cluster as the source to clone 12 DEV DB.

Not choosing D is because D -- we can not choose Serverless as master primary DB because severless is for infrequently use

upvoted 1 times

🗨️ 👤 **im_not_robot** 2 years, 3 months ago

A is fastest but B is most cost effective

D is wrong due to it is not recommend to use serverless db for production.

upvoted 2 times

🗨️ 👤 **Mintwater** 2 years, 1 month ago

Agree A

Agree your point - " not recommend to use serverless db for production"

upvoted 1 times

🗨️ 👤 **guau** 2 years, 3 months ago

Selected Answer: D

D Serverless is fastest and clone is supported

upvoted 2 times

🗨️ 👤 **teo2157** 2 years, 4 months ago

Selected Answer: D

Go for D as Aurora Serverless support cloning since June 2021

<https://aws.amazon.com/about-aws/whats-new/2021/06/amazon-aurora-serverless-v1-supports-fast-database-cloning/>

upvoted 2 times

🗨️ 👤 **lollyj** 2 years, 5 months ago

Selected Answer: A

I chose A over D because I don't believe serverless will be cheaper necessarily.

upvoted 3 times

🗨️ 👤 **awsjjj** 2 years, 7 months ago

Selected Answer: D

aurora server less now supports cloning since june 2021. Question is about cost effective. i am leaning towards D

upvoted 2 times

🗨️ 👤 **awsjjj** 2 years, 7 months ago

Although Personally I wouldn't recommend Serverless for production workload with the limitations comes with aurora serverless . A is not a wrong answer either

upvoted 1 times

🗨️ 👤 **JeanGat** 2 years, 7 months ago

Selected Answer: A

Going to go with A. Aurora Serverless is a good fit for applications that are not expected to serve traffic on a regular basis, such as development or test environments. So in this case, moving the master to serverless seems kind of backwards.

upvoted 3 times

🗨️ 👤 **niau** 2 years, 10 months ago

Selected Answer: A

A. D It s not correct. Why move to serverless?

upvoted 3 times

🗨️ 👤 **novice_expert** 3 years ago

Selected Answer: D

x A. Run the master in Amazon Aurora MySQL. Create 12 clones in VPC_TEST, and script the clones to be deleted and re-created nightly. (right answer before June 2021 as option D's serverless did not allow clone earlier)

x B. snapshot slow



x C. Replica not for testing

D. Run the master in Amazon Aurora MySQL using Aurora Serverless. Create 12 clones in VPC_TEST, and script the clones to be deleted and re-created nightly. (right answer after 6/2021, serverless v1 supports cloning to same account)

<https://aws.amazon.com/about-aws/whats-new/2021/06/amazon-aurora-serverless-v1-supports-fast-database-cloning/>

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Clone.html>

upvoted 2 times

  **awsmonster** 3 years, 4 months ago

A.

Although the question does not mention any info about the production database. I am not convinced to move the production to Aurora Serverless, with these limitations in place:

Aurora Serverless v1 doesn't support the following features:

Aurora global databases

Aurora multi-master clusters

Aurora Replicas



AWS Identity and Access Management (IAM) database authentication

Backtracking in Aurora

Database activity streams



Performance Insights

upvoted 4 times

  **guau** 2 years, 3 months ago

Severless V2 has not most of that limitations



upvoted 1 times

  **scottkerker** 3 years, 5 months ago

The best answer will be D after Jun. 21, 2021. According to <https://aws.amazon.com/about-aws/whats-new/2021/06/amazon-aurora-serverless-v1-supports-fast-database-cloning/>,



after Jun. 21, 2021, Amazon Aurora allows you to create clones between Aurora Serverless v1 and provisioned Aurora DB clusters to enable quick sharing of data, i.e., you can create Aurora Serverless v1 clones from Aurora provisioned DB clusters, and you can also create provisioned clones from Aurora Serverless v1 DB clusters.

upvoted 3 times

  **guru_ji** 3 years, 6 months ago

I got this Question in exam.

upvoted 2 times

  **guru_ji** 3 years, 7 months ago

Answer: A

upvoted 2 times

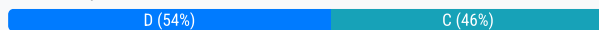
A large ecommerce company uses Amazon DynamoDB to handle the transactions on its web portal. Traffic patterns throughout the year are usually stable; however, a large event is planned. The company knows that traffic will increase by up to 10 times the normal load over the 3-day event. When sale prices are published during the event, traffic will spike rapidly.

How should a Database Specialist ensure DynamoDB can handle the increased traffic?

- A. Ensure the table is always provisioned to meet peak needs
- B. Allow burst capacity to handle the additional load
- C. Set an AWS Application Auto Scaling policy for the table to handle the increase in traffic
- D. Preprovision additional capacity for the known peaks and then reduce the capacity after the event

Suggested Answer: B

Community vote distribution



sachin Highly Voted 2 years, 11 months ago

C is the correct.

D is not correct because in Dynamodb when you scale up the capacity your data partition will increase accorssing to your RCU and WCU, but when you scale down the partition remain unchnaged, so the per table RCU and WCU will give poor performance.

I think Auto Scaling is the correct way is such situation.

upvoted 22 times

minhntm 2 years, 10 months ago

correct, I'm surprised that no one talk about it. Once you add more capacity, it's really hard to reduce

upvoted 7 times

Mintwater 2 years, 1 month ago

<https://docs.aws.amazon.com/autoscaling/application/userguide/application-auto-scaling-tutorial.html>

After completing this tutorial, you'll know how to:

Use scheduled scaling to add extra capacity to meet a heavy load before it arrives, and then remove the extra capacity when it's no longer required.

Use a target tracking scaling policy to scale your application based on current resource utilization.

Vote for C

upvoted 2 times

leotoras 1 year, 10 months ago

this document regards EC2 auto scaling, not DynamoDB scaling

upvoted 1 times

kerl 1 year, 9 months ago

<https://docs.aws.amazon.com/autoscaling/application/userguide/what-is-application-auto-scaling.html> and

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

With Application Auto Scaling, you create a scaling policy for a table or a global secondary index. The scaling policy specifies whether you want to scale read capacity or write capacity (or both), and the minimum and maximum provisioned capacity unit settings for the table or index.

Answer: C

upvoted 1 times

BillyMadison Highly Voted 3 years, 7 months ago

I'm going with D because we know about the increased traffic in advance because it will be due to a sale.

Burst capacity is fine for unknown spikes up to 5 minutes. This even is for 3 days.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-partition-key-design.html#bp-partition-key-throughput-bursting>

"DynamoDB provides some flexibility in your per-partition throughput provisioning by providing burst capacity. Whenever you're not fully using a partition's throughput, DynamoDB reserves a portion of that unused capacity for later bursts of throughput to handle usage spikes. DynamoDB currently retains up to 5 minutes (300 seconds) of unused read and write capacity. During an occasional burst of read or write activity, these extra capacity units can be consumed quickly—even faster than the per-second provisioned throughput capacity that you've defined for your table.

DynamoDB can also consume burst capacity for background maintenance and other tasks without prior notice.

Note that these burst capacity details might change in the future."

upvoted 15 times

🗳️ 👤 **MultiAZ** Most Recent 1 year, 4 months ago

Selected Answer: D

The answer is D, as we know about the event beforehand

Furthermore, C will have an issue getting 10x performance quickly enough because of the cooldown.

upvoted 1 times

🗳️ 👤 **sonu6252** 1 year, 4 months ago

D.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.html#HowItWorks.ProvisionedThroughput.M>

upvoted 1 times

🗳️ 👤 **rrshah83** 1 year, 5 months ago

Selected Answer: C

scheduled auto-scaling

upvoted 1 times

🗳️ 👤 **Santix** 1 year, 7 months ago

D is correct, because prewarm:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.html#HowItWorks.OnDemand>

upvoted 1 times

🗳️ 👤 **Germaneli** 1 year, 8 months ago

Selected Answer: C

Scheduled scaling, as one way of Application Auto Scaling, is available for DynamoDB tables and global secondary indexes. It allows to "scale a resource one time only or on a recurring schedule". I understand that this is what we need for the one-time event, and it's even automated (option D is not automated).

<https://docs.aws.amazon.com/autoscaling/application/userguide/application-auto-scaling-scheduled-scaling.html>

upvoted 1 times

🗳️ 👤 **orlvas** 1 year, 9 months ago

D

In summary, Autoscaling requires consecutive data points where the target utilization value is being breached to scale up a DynamoDB table. For this reason Autoscaling is not recommended as a solution for dealing with spiked workloads.

upvoted 1 times

🗳️ 👤 **IhorK** 1 year, 9 months ago

Selected Answer: C

Amazon DynamoDB auto scaling uses the AWS Application Auto Scaling service to dynamically adjust provisioned throughput capacity on your behalf, in response to actual traffic patterns.

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

upvoted 1 times

🗳️ 👤 **Paulv82003** 1 year, 11 months ago

Have everyone forgot cooldown during AWS Application Auto Scaling policy? We know we need to increase by 10x for AWS Application Auto Scaling policy with cooldowns it will take time to get there.

upvoted 2 times

🗳️ 👤 **sguinales** 1 year, 6 months ago

agree and in questo said "traffic will spike rapidly" autoscaling here bad performance instead provisioned, because you know when is going to be a spike and can be prepared.

upvoted 1 times

🗳️ 👤 **aviathor** 1 year, 11 months ago

Selected Answer: D

A. This is achieved by D, but D is more precise.

B. Does DynamoDB support burst capacity?

C. The question is not about the application, but about DynamoDB

D. Using provisioned capacity to meet the expected demand is one way of doing it. Using provisioned capacity with auto-scaling would also work. And of course on-demand would be an option.



upvoted 1 times

  **guau** 2 years, 3 months ago

Selected Answer: C



C- I will go with autoscaling. Why change 2 times config, when autoscaling is designed for that..

upvoted 1 times

  **leotoras** 1 year, 10 months ago

because the usage will spike rapidly, if you have pre provisioned, you dont waste time scaling



upvoted 1 times

  **im_not_robot** 2 years, 3 months ago

D is wrong because on-demand to maximum 2 times of previous peak, it can not scale to 10x.

<https://aws.amazon.com/premiumsupport/knowledge-center/on-demand-table-throttling-dynamodb/>


upvoted 1 times

  **renfdo** 2 years, 3 months ago

Selected Answer: C

C is correct. AWS always recomend to use auto scaling when you can predict usage. I now that I expect 10x more traffic.



upvoted 1 times

  **lollyj** 2 years, 5 months ago

Selected Answer: D

I'm going with D because auto scaling on the application doesn't mean the DB can accommodate the increased RW on the DB. Since the peak traffic is predictable then it may be best to pre-provision ahead and reduce after sale is over. I may be wrong though

upvoted 1 times

  **awsjjj** 2 years, 7 months ago

Selected Answer: D

I go with D. although If the answer D also includes autoscaling it would have been easy to choose D

upvoted 1 times

  **SachinGoel** 2 years, 7 months ago

Selected Answer: C

C. Set an AWS Application Auto Scaling policy for the table to handle the increase in traffic

upvoted 1 times

A Database Specialist is migrating an on-premises Microsoft SQL Server application database to Amazon RDS for PostgreSQL using AWS DMS. The application requires minimal downtime when the RDS DB instance goes live. What change should the Database Specialist make to enable the migration?

- A. Configure the on-premises application database to act as a source for an AWS DMS full load with ongoing change data capture (CDC)
- B. Configure the AWS DMS replication instance to allow both full load and ongoing change data capture (CDC)
- C. Configure the AWS DMS task to generate full logs to allow for ongoing change data capture (CDC)
- D. Configure the AWS DMS connections to allow two-way communication to allow for ongoing change data capture (CDC)

Suggested Answer: A

Reference:

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-import-data/>

Community vote distribution

A (100%)

 **awscamus** Highly Voted 3 years, 7 months ago

Sorry, I mean A. We cannot define full load or CDC in Replication instance
upvoted 10 times

 **Kamazani** 3 years, 7 months ago

B is the answer.....You define full load and CDC
when creating the replication task
upvoted 2 times

 **novice_expert** Highly Voted 3 years ago

Selected Answer: A

C&D unrelated

A. "AWS DMS full load with ongoing change data capture (CDC)" (Note WITH means 1 task)

B. allow both full load and ongoing change data capture (CDC)

(note AND, means 2 tasks)

There are two types of ongoing replication tasks:

Full load plus CDC – The task migrates existing data and then updates the target database based on changes to the source database.

CDC only – The task migrates ongoing changes after you have data on your target database.

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Task.CDC.html

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_ReplicationInstance.html

upvoted 5 times

 **aviathor** Most Recent 1 year, 11 months ago

Selected Answer: A

The answers are not very clearly stated, but A is the best option

A suggests that the source instance needs to be configured for AWS DMS full load and data capture. There is some truth to that since binary logging (transaction logs) need to be enabled.

xB suggest that it is the instance that is configured for full load and CDC, but it is really the task

xC suggests that the task should be configured to generate full logs?

xD configure two-way communication

upvoted 1 times

 **RotterDam** 3 years, 2 months ago

whoever makes these questions is probably not sober -

answer choice (A) is super apparent to the point of silliness - what ELSE would be the source?

answer choice (B) says replication instance needs to be configured to allow Full Load + CDC. The INSTANCE doesn't do this. The Task it hosts does!
Do they intend both to be the one and the same?

upvoted 4 times

🗳️ 👤 **awsmonster** 3 years, 4 months ago

Ans A: Full load and CDC are defined in a DMS Task.

upvoted 2 times

🗳️ 👤 **Dip11** 3 years, 6 months ago

Question not clear. A seems more relevant.

upvoted 4 times

🗳️ 👤 **Aesthet** 3 years, 6 months ago

"requires minimal downtime when the RDS DB instance goes live"

in order to do CDC: "you must first ensure that ARCHIVELOG MODE is on to provide information to LogMiner. AWS DMS uses LogMiner to read information from the archive logs so that AWS DMS can capture changes"

So my answer is A

upvoted 2 times

🗳️ 👤 **Zhongkai** 3 years, 6 months ago

<https://docs.aws.amazon.com/dms/latest/sbs/chap-oracle2postgresql.steps.configureoracle.html> says "If you want to capture and apply changes (CDC), then you also need the following privileges."

so A is the correct answer

upvoted 1 times

🗳️ 👤 **JobinAkaJoe** 3 years, 7 months ago

This question is bit vague. I will go with A

upvoted 1 times

🗳️ 👤 **Ashoks** 3 years, 7 months ago

A is the answer

upvoted 2 times

🗳️ 👤 **awscamus** 3 years, 7 months ago

B since Migrate existing data and replicate ongoing changes (full load + change data capture (CDC)) – To migrate data with minimal downtime, AWS DMS can migrate the existing data and replicate the data changes from the source to the target until the cutover. This migration type is best for small and medium databases that require minimal downtime, which only lasts for the duration of the cutover.

upvoted 2 times

🗳️ 👤 **MultiAZ** 3 years, 7 months ago

Answer is A. You need to to a DMS task that is "full+cdc", this has nothing to do with the DMS instance in B

upvoted 2 times

🗳️ 👤 **See111** 3 years, 7 months ago

Answer is A replication instance can't do full and cdc task .only dms do.

upvoted 2 times

🗳️ 👤 **Jiang_aws1** 2 years, 7 months ago

<https://docs.aws.amazon.com/dms/latest/sbs/chap-rdsoracle2aurora.steps.createreplicationinstance.html>

A DMS replication instance performs the actual data migration between source and target. The replication instance also caches the transaction logs during the migration. How much CPU and memory capacity a replication instance has influences the overall time required for the migration.

upvoted 1 times

🗳️ 👤 **BillyC** 3 years, 8 months ago

B i think... im not sure

upvoted 4 times

🗳️ 👤 **Mohitreddgp** 3 years, 8 months ago

A looks fine

upvoted 2 times

A financial company has allocated an Amazon RDS MariaDB DB instance with large storage capacity to accommodate migration efforts. Post-migration, the company purged unwanted data from the instance. The company now want to downsize storage to save money. The solution must have the least impact on production and near-zero downtime.


Which solution would meet these requirements?

- A. Create a snapshot of the old databases and restore the snapshot with the required storage
- B. Create a new RDS DB instance with the required storage and move the databases from the old instances to the new instance using AWS DMS
- C. Create a new database using native backup and restore
- D. Create a new read replica and make it the primary by terminating the existing primary

Suggested Answer: A

Community vote distribution

B (100%)

 **BillyMadison** Highly Voted 3 years, 8 months ago


<https://aws.amazon.com/premiumsupport/knowledge-center/rds-db-storage-size/>"After you create an Amazon RDS DB instance, you can't modify the allocated storage size of the DB instance to decrease the total storage space it uses. To decrease the storage size of your DB instance, create a new DB instance that has less provisioned storage size. Then, migrate your data into the new DB instance using one of the following methods:

Use the database engine's native dump and restore method.

Note: This method causes some downtime.

Use AWS Database Migration Service (AWS DMS) for minimal downtime."

upvoted 24 times

 **RotterDam** Highly Voted 3 years, 2 months ago


Selected Answer: B

- A. Create a snapshot of the old databases and restore the snapshot with the required storage (Downtime)
 - B. Create a new RDS DB instance with the required storage and move the databases from the old instances to the new instance using AWS DMS (no downtime)
 - C. Create a new database using native backup and restore (Downtime like A)
 - D. Create a new read replica and make it the primary by terminating the existing primary (wont change the storage)
- upvoted 5 times

 **sachin** Most Recent 2 years, 11 months ago

Its not D because For replication to operate effectively, each Read Replica should have the same amount of compute & storage resources as the source DB instance.

upvoted 4 times

 **Dantas** 3 years, 2 months ago

Selected Answer: B

B. When you create a read replica you cannot set the storage size... so D is incorrect.

upvoted 2 times

 **yahoos** 3 years, 3 months ago

B

Almost no downtime > AWS DMS

upvoted 1 times

 **aws_enthu** 3 years, 4 months ago


Isn't A?

upvoted 1 times

 **user0001** 3 years, 2 months ago

no, this will not shrink the size

upvoted 2 times

 **Nice_Guy** 2 years, 2 months ago

key: 'near-zero downtime'.

upvoted 1 times

🗋️ 👤 **shuraosipov** 3 years, 5 months ago

Selected Answer: B

Answer is B.

You cannot downsize the database storage size on RDS.

upvoted 2 times

🗋️ 👤 **jove** 3 years, 6 months ago

Almost no downtime > AWS DMS : Option B

upvoted 2 times

🗋️ 👤 **aws4myself** 3 years, 6 months ago

D is not correct because, You cannot set the size of the read replica, which is always same as the master.

upvoted 3 times

🗋️ 👤 **stevewuoisiro** 3 years, 7 months ago

B is correct. D is close, however, you don't make the read replica the primary by terminating the existing primary.

"Create a new read replica and make it the primary by terminating the existing primary"

upvoted 1 times

🗋️ 👤 **Dantas** 3 years, 2 months ago

Yes, you can make the read replica the primary just by terminating the current primary. Anyways option D is incorrect once the replica storage will be the same as the primary.

upvoted 1 times

🗋️ 👤 **Dip11** 3 years, 7 months ago

D can be an option too. Read replicas can have different storage size and Read replicas will have less impact on production as compared to DMS.

upvoted 1 times

🗋️ 👤 **Aesthet** 3 years, 7 months ago

B looks correct according to BillyMadaison's link. Not my initial answer though

upvoted 1 times

🗋️ 👤 **myutran** 3 years, 7 months ago

Ans: B

upvoted 1 times

🗋️ 👤 **JobinAkaJoe** 3 years, 7 months ago

B looks to be the correct answer

upvoted 1 times

🗋️ 👤 **Ashoks** 3 years, 7 months ago

Yes. B is the answer

snapshot will retain same size. DB needs to be recreated and DMS is for near zero

upvoted 3 times

🗋️ 👤 **halol** 3 years, 8 months ago

B is corect

upvoted 1 times

🗋️ 👤 **Ebi** 3 years, 8 months ago

B is correct

upvoted 1 times

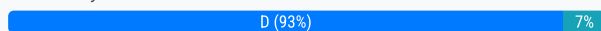
A large financial services company requires that all data be encrypted in transit. A Developer is attempting to connect to an Amazon RDS DB instance using the company VPC for the first time with credentials provided by a Database Specialist. Other members of the Development team can connect, but this user is consistently receiving an error indicating a communications link failure. The Developer asked the Database Specialist to reset the password a number of times, but the error persists.

Which step should be taken to troubleshoot this issue?

- A. Ensure that the database option group for the RDS DB instance allows ingress from the Developer machine's IP address
- B. Ensure that the RDS DB instance's subnet group includes a public subnet to allow the Developer to connect
- C. Ensure that the RDS DB instance has not reached its maximum connections limit
- D. Ensure that the connection is using SSL and is addressing the port where the RDS DB instance is listening for encrypted connections

Suggested Answer: B

Community vote distribution



k115 Highly Voted 3 years, 8 months ago

D is correct

upvoted 8 times

Germaneli Most Recent 1 year, 8 months ago

Selected Answer: D

It's noted in the question that the connection needs to be encrypted.

It's also noted that the logon credentials are confirmed, they should be right.

B is out of question, there is no such thing as a subnet group to control access.

D is the only reasonable option.

upvoted 1 times

Pranava_GCP 1 year, 8 months ago

Selected Answer: D

D. Ensure that the connection is using SSL and is addressing the port where the RDS DB instance is listening for encrypted connections

upvoted 2 times

rInd2000 2 years, 11 months ago

Selected Answer: D

A lot of time this issue has happened to me, a Developer call me because is unable to connect to Redshift or any other DB and the problem is that they forgot to check "use SSL" in DBeaver. :)

I will go with D:

upvoted 4 times

ryuhei 2 years, 11 months ago

Selected Answer: D

Ofcourse it's decided to be D

upvoted 1 times

awsguys 3 years ago

d. as all data in transit be encrypted

upvoted 2 times

novice_expert 3 years ago

Selected Answer: D

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/SQLServer.Concepts.General.SSL.Using.html>

D. addressing the port where the RDS DB instance is listening for encrypted connections (communications connection failure => port)

upvoted 2 times

kret 3 years, 1 month ago

Selected Answer: D

"Other members of the Development team are able to connect", so obviously not B

upvoted 2 times

🗳️ 👤 **soyyodario** 3 years, 4 months ago

Selected Answer: B

"firm mandates that all data in transit be encrypted" but "Other members of the Development team are able to connect", then SSL connections are ok.

B is the answer that has more sense for communication error.

upvoted 1 times

🗳️ 👤 **Shunpin** 3 years, 5 months ago

Selected Answer: D

"a communications connection failure". It's about communication failure during connection.

upvoted 1 times

🗳️ 👤 **gelsm** 3 years, 6 months ago

Answer D: <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/SQLServer.Concepts.General.SSL.Using.html>

upvoted 1 times

🗳️ 👤 **Dip11** 3 years, 6 months ago

D for sure.

upvoted 1 times

🗳️ 👤 **shantest1** 3 years, 6 months ago

D is correct

upvoted 1 times

🗳️ 👤 **myutran** 3 years, 6 months ago

Ans: D

upvoted 2 times

🗳️ 👤 **JobinAkaJoe** 3 years, 6 months ago

D is my choice

upvoted 2 times

🗳️ 👤 **Ashoks** 3 years, 7 months ago

D is the answer

upvoted 2 times

🗳️ 👤 **Ebi** 3 years, 7 months ago

D is answer

upvoted 1 times

A company is running Amazon RDS for MySQL for its workloads. There is downtime when AWS operating system patches are applied during the Amazon RDS- specified maintenance window.


What is the MOST cost-effective action that should be taken to avoid downtime?

- A. Migrate the workloads from Amazon RDS for MySQL to Amazon DynamoDB
- B. Enable cross-Region read replicas and direct read traffic to them when Amazon RDS is down
- C. Enable a read replica and direct read traffic to it when Amazon RDS is down
- D. Enable an Amazon RDS for MySQL Multi-AZ configuration

Suggested Answer: C

Community vote distribution

D (100%)

 **BillyMadison** Highly Voted 3 years, 8 months ago

Going with D for now

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-required-maintenance/>

To minimize downtime, modify the Amazon RDS DB instance to a Multi-AZ deployment. For Multi-AZ deployments, OS maintenance is applied to the secondary instance first, then the instance fails over, and then the primary instance is updated. The downtime is during failover. For more information, see Maintenance for Multi-AZ Deployments.

<https://aws.amazon.com/rds/faqs/>

The availability benefits of Multi-AZ also extend to planned maintenance. For example, with automated backups, I/O activity is no longer suspended on your primary during your preferred backup window, since backups are taken from the standby. In the case of patching or DB instance class scaling, these operations occur first on the standby, prior to automatic fail over. As a result, your availability impact is limited to the time required for automatic failover to complete.


upvoted 14 times

 **pan24** Highly Voted 3 years, 8 months ago

Ans: D

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-required-maintenance/>

upvoted 5 times

 **Germaneli** Most Recent 1 year, 8 months ago

Selected Answer: D

A is a diverter.

B + C are only for reading, that doesn't help a Production database to avoid downtime.

D is documented to help reduce downtime during OS patch cycles.

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-required-maintenance/>

upvoted 1 times

 **Pranava_GCP** 1 year, 8 months ago

Selected Answer: D

D. Enable an Amazon RDS for MySQL Multi-AZ configuration

upvoted 1 times


 **IhorK** 1 year, 9 months ago

Selected Answer: D

For OS maintenance ("AWS operating system patches are applied"), OS maintenance is applied to the secondary instance first, then the instance fails over, and then the primary instance is updated.

<https://repost.aws/knowledge-center/rds-required-maintenance>

upvoted 1 times

 **adelcold** 1 year, 11 months ago

Selected Answer: D

<https://repost.aws/knowledge-center/rds-required-maintenance>

upvoted 1 times

🗨️ 👤 **praffuln** 3 years ago

Selected Answer: D

Obviously D is correct.

upvoted 2 times

🗨️ 👤 **novice_expert** 3 years ago

Selected Answer: D

C is the answer if the workload is read-only.

D is the answer to do the maintenance for R+W workload with reduced outage

upvoted 1 times

🗨️ 👤 **stevewuoisiro** 3 years, 6 months ago

D is correct: "Single-AZ deployments are unavailable for a few minutes. Multi-AZ deployments are unavailable for the time it takes the instance to failover (usually about 60 seconds) if the Availability Zone is affected by the maintenance. If only the secondary Availability Zone is affected, then there is no failover or downtime. "

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-required-maintenance/>

upvoted 1 times

🗨️ 👤 **Dip11** 3 years, 6 months ago

D for sure.

upvoted 1 times

🗨️ 👤 **shantest1** 3 years, 7 months ago

Ans: D

upvoted 2 times

🗨️ 👤 **myutran** 3 years, 7 months ago

-Ans: D

upvoted 2 times

🗨️ 👤 **JobinAkaJoe** 3 years, 7 months ago

A. Migrate the workloads from Amazon RDS for MySQL to Amazon DynamoDB

DynamoDB is not appropriate to take over RDS workload.

B. Enable cross-Region read replicas and direct read traffic to then when Amazon RDS is down

What about writes ? why do you need cross-region when you can have read-replicas in same region

C. Enable a read replicas and direct read traffic to it when Amazon RDS is down

What about write workload ?

D. Enable an Amazon RDS for MySQL Multi-AZ configuration

Helps in reducing the outage required for the maintenance(just a failover)

C is the answer if the workload is read-only.

D is the answer to do the maintenance for R+W workload with reduced outage

upvoted 4 times

🗨️ 👤 **Ashoks** 3 years, 7 months ago

Yes it is D

upvoted 2 times

🗨️ 👤 **Billhardy** 3 years, 7 months ago

Ans is D

upvoted 1 times

🗨️ 👤 **Ebi** 3 years, 7 months ago

Answer is D

upvoted 1 times

🗨️ 👤 **[Removed]** 3 years, 7 months ago

D. 100%

upvoted 2 times

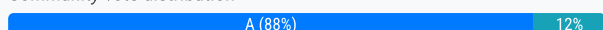
A Database Specialist must create a read replica to isolate read-only queries for an Amazon RDS for MySQL DB instance. Immediately after creating the read replica, users that query it report slow response times.

What could be causing these slow response times?

- A. New volumes created from snapshots load lazily in the background
- B. Long-running statements on the master
- C. Insufficient resources on the master
- D. Overload of a single replication thread by excessive writes on the master

Suggested Answer: B

Community vote distribution



pan24 Highly Voted 3 years, 8 months ago

ANS: A

snapshot is lazy loaded If the volume is accessed where the data is not loaded, the application accessing the volume encounters a higher latency than normal while the data gets loaded

upvoted 15 times

gelsm 3 years, 6 months ago

<https://aws.amazon.com/about-aws/whats-new/2019/11/amazon-ebs-fast-snapshot-restore-eliminates-need-for-prewarming-data-into-volumes-created-snapshots/>

upvoted 1 times

hariti_crafting 2 years, 10 months ago

question is about RDS this link is irrelevant not need to confuse ebs and rds

upvoted 4 times

RotterDam Highly Voted 3 years, 2 months ago

Selected Answer: A

Unlike Aurora - RDS does not have a warm pool of cache (unless its POSTGRES and you are using CCM there - but there is MYSQL). First touch penalty. Can be mitigated with doing a select * from all tables

upvoted 10 times

aviathor Most Recent 1 year, 11 months ago

Selected Answer: B

An active, long-running transaction can slow the process of creating the read replica. We recommend that you wait for long-running transactions to complete before creating a read replica. If you create multiple read replicas in parallel from the same source DB instance, Amazon RDS takes only one snapshot at the start of the first create action.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html#USER_ReadRepl.Create

upvoted 2 times

Rayls2000 2 years, 10 months ago

Ans: B

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-mysql-high-replica-lag/>

If the replica SQL_THREAD is the source of replication delays, then those delays could be caused by the following:

- Long-running queries on the primary DB instance
 - Insufficient DB instance class size or storage
 - Parallel queries run on the primary DB instance
 - Binary logs synced to the disk on the replica DB instance
 - Binlog_format on the replica is set to ROW
 - Replica creation lag
- upvoted 2 times

🗨️ 👤 **ray6li** 2 years, 10 months ago

Correction: Ans: A. Lag != slow response time.

upvoted 3 times

🗨️ 👤 **novice_expert** 3 years ago

Selected Answer: A

<https://aws.amazon.com/about-aws/whats-new/2019/11/amazon-ebs-fast-snapshot-restore-eliminates-need-for-prewarming-data-into-volumes-created-snapshots/>

upvoted 3 times

🗨️ 👤 **hariti_crafting** 2 years, 10 months ago

question is about RDS this link is irrelevant not need to confuse ebs and rds

upvoted 4 times

🗨️ 👤 **GMartinelli** 3 years, 6 months ago

Selected Answer: A

Option A.

upvoted 2 times

🗨️ 👤 **shantest1** 3 years, 6 months ago

Ans: A

upvoted 1 times

🗨️ 👤 **JobinAkaJoe** 3 years, 7 months ago

A is my choice

upvoted 1 times

🗨️ 👤 **Ashoks** 3 years, 7 months ago

Yes, it is A

upvoted 2 times

🗨️ 👤 **Ebi** 3 years, 7 months ago

A is correct

upvoted 2 times

🗨️ 👤 **firbhat** 3 years, 7 months ago

Ans A

When you spin up a new replica, its EBS volume loads lazily in the background

upvoted 2 times

🗨️ 👤 **BillyC** 3 years, 7 months ago

Yes A is Correct

upvoted 3 times

A company developed an AWS CloudFormation template used to create all new Amazon DynamoDB tables in its AWS account. The template configures provisioned throughput capacity using hard-coded values. The company wants to change the template so that the tables it creates in the future have independently configurable read and write capacity units assigned.

Which solution will enable this change?

- A. Add values for the rcuCount and wcuCount parameters to the Mappings section of the template. Configure DynamoDB to provision throughput capacity using the stack's mappings.
- B. Add values for two Number parameters, rcuCount and wcuCount, to the template. Replace the hard-coded values with calls to the Ref intrinsic function, referencing the new parameters.
- C. Add values for the rcuCount and wcuCount parameters as outputs of the template. Configure DynamoDB to provision throughput capacity using the stack outputs.
- D. Add values for the rcuCount and wcuCount parameters to the Mappings section of the template. Replace the hard-coded values with calls to the Ref intrinsic function, referencing the new parameters.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ **sk1974** 2 years, 2 months ago

I will say 'B' too . the key for me to decide is the word 'number' since he wcu and rcu are numeric values. Not sure if my interpretation is correct ,but , will go with B

upvoted 1 times

🗳️ **rlnd2000** 2 years, 11 months ago

Selected Answer: B

A and D are out, parameters are useful when you know the value for me mapping are like constant in any program language we can use them for Region, Accounts, AZ, etc and the question says "...allocate independently variable read and write capacity..." the word variable is key for me here.

C- is out Output is for importing values into other stack.

B is correct Use the optional Parameters section to customize your templates.

Parameters enable you to input custom values to your template each time you create or update a stack. from

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/parameters-section-structure.html>

D -

upvoted 3 times

🗳️ **Radhaghosh** 3 years ago

Selected Answer: B

"The organization want to modify the template in order to allocate independently variable read and write capacity units to future tables." this is the vital sentence in the question. Only Option is Parameter. Option B

upvoted 1 times

🗳️ **novice_expert** 3 years ago

Selected Answer: B

Add values for two Number parameters, rcuCount and wcuCount, to the template. Replace the hard-coded values with calls to the Ref intrinsic function, referencing the new parameters.

upvoted 1 times

🗳️ **GMartinelli** 3 years, 5 months ago

Selected Answer: B

Option B

upvoted 2 times

🗳️ 👤 **guru_ji** 3 years, 6 months ago

I got this Question in exam.

upvoted 1 times

🗳️ 👤 **umatrilok** 3 years, 7 months ago

definetly A

upvoted 1 times

🗳️ 👤 **Aesthet** 3 years, 7 months ago

probably B

upvoted 1 times

🗳️ 👤 **Windy** 3 years, 7 months ago

Why not A? It is obvious that there are multiple tables and rcu and wcu should be defined for each of them. I think mapping is a must here.

upvoted 3 times

🗳️ 👤 **Zhongkai** 3 years, 7 months ago

Mapping is not used like this way. Check <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/mappings-section-structure.html>.

We should use Parameter here so that the user could input what RCU/WCU they need when deploying a new stack.

upvoted 1 times

🗳️ 👤 **Aesthet** 3 years, 7 months ago

Input parameter and FindInMap

You can use an input parameter with the Fn::FindInMap function to refer to a specific value in a map. For example, suppose you have a list of regions and environment types that map to a specific AMI ID. You can select the AMI ID that your stack uses by using an input parameter (EnvironmentType). To determine the region, use the AWS::Region pseudo parameter, which gets the AWS Region in which you create the stack.

I am not saying that B is not right answer though

upvoted 1 times

🗳️ 👤 **ImprovMAN** 2 years, 5 months ago

You can't include parameters, pseudo parameters, or intrinsic functions in the Mappings section.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/mappings-section-structure.html>

upvoted 1 times

🗳️ 👤 **myutran** 3 years, 7 months ago

Ans: B

upvoted 3 times

🗳️ 👤 **JobinAkaJoe** 3 years, 7 months ago

B is my choice

upvoted 1 times

🗳️ 👤 **Ashoks** 3 years, 7 months ago

yes B for dynamic and not pre-determined values

upvoted 3 times

🗳️ 👤 **BillyMadison** 3 years, 7 months ago

GOing with B

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/parameters-section-structure.html>

upvoted 4 times

🗳️ 👤 **Ebi** 3 years, 7 months ago

B is correct

upvoted 1 times

🗳️ 👤 **Kitty0403** 3 years, 8 months ago

Answer is B

upvoted 1 times

🗳️ 👤 **BillyC** 3 years, 8 months ago

B or C, please Thoughts

upvoted 1 times

A retail company with its main office in New York and another office in Tokyo plans to build a database solution on AWS. The company's main workload consists of a mission-critical application that updates its application data in a data store. The team at the Tokyo office is building dashboards with complex analytical queries using the application data. The dashboards will be used to make buying decisions, so they need to have access to the application data in less than 1 second.

Which solution meets these requirements?

- A. Use an Amazon RDS DB instance deployed in the us-east-1 Region with a read replica instance in the ap-northeast-1 Region. Create an Amazon ElastiCache cluster in the ap-northeast-1 Region to cache application data from the replica to generate the dashboards.
- B. Use an Amazon DynamoDB global table in the us-east-1 Region with replication into the ap-northeast-1 Region. Use Amazon QuickSight for displaying dashboard results.
- C. Use an Amazon RDS for MySQL DB instance deployed in the us-east-1 Region with a read replica instance in the ap-northeast-1 Region. Have the dashboard application read from the read replica.
- D. Use an Amazon Aurora global database. Deploy the writer instance in the us-east-1 Region and the replica in the ap-northeast-1 Region. Have the dashboard application read from the replica ap-northeast-1 Region.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ **IhorK** 1 year, 9 months ago

Selected Answer: D

- complex analytical queries
- less than 1 second
- 2 Regions

Amazon Aurora Global Database cross-Region replication latencies below 1 second, recovery point objective (RPO) of 1 second and a recovery time objective (RTO) of less than 1 minute.

upvoted 1 times

🗳️ **renfdo** 2 years, 3 months ago

Selected Answer: D

Use an Amazon Aurora global database.

upvoted 1 times

🗳️ **Arun32** 2 years, 6 months ago

Surely D :

Global Database uses storage-based replication with typical latency of less than 1 second, using dedicated infrastructure that leaves your database fully available to serve application workloads.

<https://aws.amazon.com/rds/aurora/global-database/>

upvoted 4 times

🗳️ **novice_expert** 3 years ago

Selected Answer: D

Use an Amazon Aurora global database.

upvoted 1 times

🗳️ **guru_ji** 3 years, 6 months ago

Correct Answer: D

Dynamo DB ==>> No direct analytical queries (No joins).

upvoted 2 times

🗳️ **LB** 3 years, 6 months ago

D - Aurora Global Database - The solution needs to support complex analytical queries. Which eliminates Dynamodb from the equation.

upvoted 2 times

🗳️ **sbhujbal** 3 years, 7 months ago

option B is not possible since Quicksight cannot talk with DynamoDB and hence only possible option is D
upvoted 3 times

🗨️ 👤 **Aesthet** 3 years, 7 months ago

D final answer

upvoted 3 times

🗨️ 👤 **myutran** 3 years, 7 months ago

Ans: D

upvoted 3 times

🗨️ 👤 **halol** 3 years, 7 months ago

why it's not A, I can see A and D make sense !

upvoted 1 times

🗨️ 👤 **aws4myself** 3 years, 7 months ago

Because they have asked for 1 sec replication lag.

upvoted 1 times

🗨️ 👤 **Ashoks** 3 years, 7 months ago

I would go with D. Covers cross region 1 sec latency and complex query

upvoted 4 times

🗨️ 👤 **pdboi3355** 3 years, 7 months ago

D - "This means that committed transactional changes from the writer are replicated globally to the Regions that you select, typically within 1 second." -- based on link below

upvoted 2 times

🗨️ 👤 **goodh32** 3 years, 7 months ago

Answer is B

Complex query and analytics is key to use RDS, not DynamoDB

upvoted 2 times

🗨️ 👤 **goodh32** 3 years, 7 months ago

**editing:

Answer is for Aurora (D) - which is Relational ready for complex queries

upvoted 2 times

🗨️ 👤 **qwertyuio** 3 years, 7 months ago

D. Quicksight don't support Dynamodb

upvoted 3 times

🗨️ 👤 **GeeBeeEI** 3 years, 7 months ago

I dont agree with you, see <https://aws.amazon.com/blogs/database/how-to-perform-advanced-analytics-and-build-visualizations-of-your-amazon-dynamodb-data-by-using-amazon-athena/>

upvoted 2 times

🗨️ 👤 **jove** 3 years, 5 months ago

That blog suggests to export data from DynamoDB to S3 first and use Athena to read from S3. This process is slow. In this question, the dashboard requires an instant access to the application's data.

Answer is D

upvoted 1 times

🗨️ 👤 **Ebi** 3 years, 7 months ago

Answer is D

upvoted 1 times

🗨️ 👤 **Kitty0403** 3 years, 8 months ago

Answer is B

upvoted 1 times

🗨️ 👤 **Kitty0403** 3 years, 8 months ago

Changing to B

<https://aws.amazon.com/blogs/database/aurora-postgresql-disaster-recovery-solutions-using-amazon-aurora-global-database/>

upvoted 2 times

  **Kitty0403** 3 years, 7 months ago



Sorry I meant D

upvoted 2 times

  **Kitty0403** 3 years, 7 months ago

1 sec is the key.. dyanmodb takes 2000ms for replication

upvoted 2 times

  **steves** 3 years, 8 months ago

Going with B, less than 1 second replication required.

upvoted 1 times

A company is using Amazon RDS for PostgreSQL. The Security team wants all database connection requests to be logged and retained for 180 days. The RDS for PostgreSQL DB instance is currently using the default parameter group. A Database Specialist has identified that setting the `log_connections` parameter to 1 will enable connections logging.

Which combination of steps should the Database Specialist take to meet the logging and retention requirements? (Choose two.)

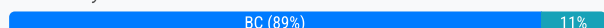
- A. Update the `log_connections` parameter in the default parameter group
- B. Create a custom parameter group, update the `log_connections` parameter, and associate the parameter with the DB instance
- C. Enable publishing of database engine logs to Amazon CloudWatch Logs and set the event expiration to 180 days
- D. Enable publishing of database engine logs to an Amazon S3 bucket and set the lifecycle policy to 180 days
- E. Connect to the RDS PostgreSQL host and update the `log_connections` parameter in the `postgresql.conf` file

Suggested Answer: AE

Reference:

<https://aws.amazon.com/blogs/database/working-with-rds-and-aurora-postgresql-logs-part-1/>

Community vote distribution



BillyMadison Highly Voted 3 years, 7 months ago

B&C

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_LogAccess.Concepts.PostgreSQL.html#USER_LogAccess.PostgreSQL.log_retention_period
To set the retention period for system logs, use the `rds.log_retention_period` parameter. You can find `rds.log_retention_period` in the DB parameter group associated with your DB instance. The unit for this parameter is minutes. For example, a setting of 1,440 retains logs for one day. The default value is 4,320 (three days). The maximum value is 10,080 (seven days). Your instance must have enough allocated storage to contain the retained log files.

To retain older logs, publish them to Amazon CloudWatch Logs. For more information, see [Publishing PostgreSQL Logs to CloudWatch Logs](#).

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_LogAccess.Concepts.PostgreSQL.html#USER_LogAccess.PostgreSQL.PublishtoCloudWatchLogs
upvoted 12 times

Pranava_GCP Most Recent 1 year, 8 months ago

Selected Answer: BC

BC are correct.

A is wrong, can not change default param group

E is wrong, can not update `postgresql.conf` file on RDS, maybe can do on EC2 ?

upvoted 1 times

ken_test1234 2 years, 1 month ago

Selected Answer: BC

you can not alter default parameter group , i read the documentation of rds and there is no integration to export logs to s3

upvoted 1 times

backbencher2022 2 years, 2 months ago

Selected Answer: BC

Clearly B&C as per this AWS URL - <https://aws.amazon.com/premiumsupport/knowledge-center/track-failed-login-rds-postgresql/>

upvoted 1 times

sk1974 2 years, 2 months ago

Why not B&D . There is no need to analyze the logs in Cloudwatch.They have to be retained in 180 days. Isn't storing data in S3 cheaper than storing them in Cloudwatch ?

upvoted 1 times

ninjalight25 2 years, 2 months ago

Selected Answer: AC

To meet the logging and retention requirements, the following steps should be taken:

Update the `log_connections` parameter in the default parameter group to enable connection logging.

Enable publishing of database engine logs to Amazon CloudWatch Logs and set the event expiration to 180 days. This will allow the logging of all database connection requests to be retained for 180 days.

upvoted 1 times

🗳️ 👤 **subzzzero** 2 years, 7 months ago

why A,E? this is misleading there is no way you can edit postgresql.conf in RDS... correct is B&C

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: BC

B. Create a custom parameter group, update the log_connections parameter, and associate the parameter with the DB instance

C. Enable publishing of database engine logs to Amazon CloudWatch Logs and set the event expiration to 180 days

upvoted 2 times

🗳️ 👤 **soyyodario** 3 years, 4 months ago

Selected Answer: BC

B and C.

A wrong:

Parameter groups

Each Amazon RDS PostgreSQL instance is associated with a parameter group that contains the engine specific configurations. The engine configurations also include several parameters that control PostgreSQL logging behavior. AWS provides the parameter groups with default configuration settings to use for your instances. However, to change the default settings, you must create a clone of the default parameter group, modify it, and attach it to your instance.

To set logging parameters for a DB instance, set the parameters in a DB parameter group and associate that parameter group with the DB instance. For more information, see Working with DB parameter groups.

upvoted 3 times

🗳️ 👤 **Dip11** 3 years, 6 months ago

B and C for sure.

upvoted 1 times

🗳️ 👤 **Aesthet** 3 years, 6 months ago

BC final answer

upvoted 2 times

🗳️ 👤 **myutran** 3 years, 6 months ago

Ans: BC

upvoted 4 times

🗳️ 👤 **JobinAkaJoe** 3 years, 6 months ago

B and C

upvoted 1 times

🗳️ 👤 **Ashoks** 3 years, 7 months ago

Yes B&C

upvoted 2 times

🗳️ 👤 **Ebi** 3 years, 7 months ago

BC correct answer

upvoted 1 times

🗳️ 👤 **Kitty0403** 3 years, 7 months ago

Answer is B and C

upvoted 1 times

🗳️ 👤 **SaulGoodman** 3 years, 7 months ago

agreed with B&C

export logs to Cloudwatch Logs

upvoted 2 times

A Database Specialist is creating a new Amazon Neptune DB cluster, and is attempting to load data from Amazon S3 into the Neptune DB cluster using the Neptune bulk loader API. The Database Specialist receives the following error:

```
`Unable to connect to s3 endpoint. Provided source = s3://mybucket/graphdata/ and region = us-east-1. Please verify your S3 configuration.`
```

Which combination of actions should the Database Specialist take to troubleshoot the problem? (Choose two.)

- A. Check that Amazon S3 has an IAM role granting read access to Neptune
- B. Check that an Amazon S3 VPC endpoint exists
- C. Check that a Neptune VPC endpoint exists
- D. Check that Amazon EC2 has an IAM role granting read access to Amazon S3
- E. Check that Neptune has an IAM role granting read access to Amazon S3

Suggested Answer: BD

Reference:

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-could-not-connect-endpoint-url/>


Community vote distribution

BE (82%)

AB (18%)

 **k115** Highly Voted 3 years, 8 months ago

B & E <https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load-tutorial-IAM.html>
upvoted 16 times

 **BillyMadison** Highly Voted 3 years, 7 months ago


BE

<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load-data.html>

"An IAM role for the Neptune DB instance to assume that has an IAM policy that allows access to the data files in the S3 bucket. The policy must grant Read and List permissions."

"An Amazon S3 VPC endpoint. For more information, see the Creating an Amazon S3 VPC Endpoint section."

upvoted 8 times

 **Germaneli** Most Recent 1 year, 8 months ago

Selected Answer: AB

- A. Neptune wants to read from S3, so S3 needs to provide read access.
- B. Yes, we can make use of an Amazon S3 VPC endpoint to access data.
- x C. There's nothing like a Neptune VPC endpoint.
- x D. We're not dealing with EC2 here.
- x E. Neptune wants to read from S3, not the other way around.

upvoted 2 times


 **Pranava_GCP** 1 year, 8 months ago

Selected Answer: BE

- B. Check that an Amazon S3 VPC endpoint exists
- E. Check that Neptune has an IAM role granting read access to Amazon S3

<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load-data.html>

upvoted 1 times

 **lollyj** 2 years, 5 months ago

Selected Answer: BE

My final answer

upvoted 2 times

 **novice_expert** 3 years ago

Selected Answer: BE

<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load-data.html>

B. Check that an Amazon S3 VPC endpoint exists

E. Check that Neptune has an IAM role granting read access to Amazon S3

An IAM role for the Neptune DB instance to assume that has an IAM policy that allows access to the data files in the S3 bucket. The policy must grant Read and List permissions.

upvoted 1 times

  **RotterDam** 3 years, 2 months ago

Selected Answer: BE

BE are the correct answers

- S3 VPC endpoint has to exist

- Neptune must have the IAM role and policies with access to S3 bucket

upvoted 3 times



  **soyyodario** 3 years, 4 months ago

Selected Answer: BE

B and E

<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load-data.html>

upvoted 2 times

  **Aesthet** 3 years, 6 months ago

BE final answer

upvoted 1 times

  **manan728** 3 years, 6 months ago



Guys this question was asked in my exam that I passed. B and E are correct.

upvoted 4 times

  **myutran** 3 years, 6 months ago



Ans: BE

upvoted 2 times

  **Umer24** 3 years, 7 months ago

send me a note ryan23680 at yahoo for new aws database questions to find the correct answers.

upvoted 1 times

  **Umer24** 3 years, 7 months ago

Question-121

A company's database specialist disabled TLS on an Amazon DocumentDB cluster to perform benchmarking tests. A few days after this change was implemented, a database specialist trainee accidentally deleted multiple tables. The database specialist restored the database from available snapshots. An hour after restoring the cluster, the database specialist is still unable to connect to the new cluster endpoint. What should the database specialist do to connect to the new, restored Amazon DocumentDB cluster?


A. Change the restored cluster's parameter group to the original cluster's custom parameter group.

B. Change the restored cluster's parameter group to the Amazon DocumentDB default parameter group.

C. Configure the interface VPC endpoint and associate the new Amazon DocumentDB cluster.

D. Run the syncInstances command in AWS DataSync.

upvoted 1 times

  **Glendon** 3 years, 7 months ago



What is the answer for this? A?

upvoted 2 times

  **JobinAkaJoe** 3 years, 7 months ago



B and E my choice

upvoted 1 times

  **Ashoks** 3 years, 7 months ago



Ans is B&E

upvoted 2 times

  **Ebi** 3 years, 7 months ago

Answer is BE

upvoted 2 times

  **BillyC** 3 years, 7 months ago

B and E here

upvoted 2 times

A database specialist manages a critical Amazon RDS for MySQL DB instance for a company. The data stored daily could vary from .01% to 10% of the current database size. The database specialist needs to ensure that the DB instance storage grows as needed. What is the MOST operationally efficient and cost-effective solution?

- A. Configure RDS Storage Auto Scaling.
- B. Configure RDS instance Auto Scaling.
- C. Modify the DB instance allocated storage to meet the forecasted requirements.
- D. Monitor the Amazon CloudWatch FreeStorageSpace metric daily and add storage as required.

Suggested Answer: B

Community vote distribution

A (100%)

🗳️ 👤 **ken_test1234** 2 years, 1 month ago

Selected Answer: A

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.StorageTypes.html#USER_PIOPS.Autoscaling
upvoted 2 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: A

Storage auto scaling
upvoted 2 times

🗳️ 👤 **Dantas** 3 years, 2 months ago

Selected Answer: A

There's no Amazon RDS for MySQL DB instance autoscaling service (it's only available for Aurora). And it wouldn't solve the storage issue.
upvoted 2 times

🗳️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: A

agree with other comments
upvoted 2 times

🗳️ 👤 **peacegrace** 3 years, 4 months ago

Selected Answer: A

Cost Effective
upvoted 2 times

🗳️ 👤 **Shunpin** 3 years, 5 months ago

Ans: A

From the AWS Console, you can only see storage section with autoscaling.
upvoted 3 times

🗳️ 👤 **shuraosipov** 3 years, 5 months ago

Selected Answer: A

Answer is A.

If your workload is unpredictable, you can enable storage autoscaling for an Amazon RDS DB instance. With storage autoscaling enabled, when Amazon RDS detects that you are running out of free database space it automatically scales up your storage.

<https://aws.amazon.com/about-aws/whats-new/2019/06/rds-storage-auto-scaling/>

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.StorageTypes.html#USER_PIOPS.Autoscaling

upvoted 4 times

🗳️ 👤 **gelsm** 3 years, 6 months ago

A. Configure RDS Storage Auto Scaling.

<https://aws.amazon.com/vi/about-aws/whats-new/2019/06/rds-storage-auto-scaling/>

"operationally efficient and cost-effective"

RDS Storage Auto Scaling automatically scales storage capacity in response to growing database workloads, with zero downtime.

There is no additional cost for RDS Storage Auto Scaling.

upvoted 4 times

🗲️ 👤 **Dip11** 3 years, 6 months ago

A is correct ans.

upvoted 1 times

🗲️ 👤 **yossanjp** 3 years, 7 months ago

Ans.A

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.StorageTypes.html#USER_PIOPS.Autoscaling

upvoted 1 times

🗲️ 👤 **Aesthet** 3 years, 7 months ago

A final answer

upvoted 2 times

🗲️ 👤 **thanhphan** 3 years, 7 months ago

I'll go with A.

<https://aws.amazon.com/vi/about-aws/whats-new/2019/06/rds-storage-auto-scaling/>

upvoted 1 times

🗲️ 👤 **shantest1** 3 years, 8 months ago

A. Answer

upvoted 3 times

A company is due for renewing its database license. The company wants to migrate its 80 TB transactional database system from on-premises to the AWS Cloud.

The migration should incur the least possible downtime on the downstream database applications. The company's network infrastructure has limited network bandwidth that is shared with other applications.

Which solution should a database specialist use for a timely migration?

- A. Perform a full backup of the source database to AWS Snowball Edge appliances and ship them to be loaded to Amazon S3. Use AWS DMS to migrate change data capture (CDC) data from the source database to Amazon S3. Use a second AWS DMS task to migrate all the S3 data to the target database.
- B. Perform a full backup of the source database to AWS Snowball Edge appliances and ship them to be loaded to Amazon S3. Periodically perform incremental backups of the source database to be shipped in another Snowball Edge appliance to handle syncing change data capture (CDC) data from the source to the target database.
- C. Use AWS DMS to migrate the full load of the source database over a VPN tunnel using the internet for its primary connection. Allow AWS DMS to handle syncing change data capture (CDC) data from the source to the target database.
- D. Use the AWS Schema Conversion Tool (AWS SCT) to migrate the full load of the source database over a VPN tunnel using the internet for its primary connection. Allow AWS SCT to handle syncing change data capture (CDC) data from the source to the target database.

Suggested Answer: D

Community vote distribution

A (82%)

B (18%)

🗳️ 👤 **frankzeng** Highly Voted 👍 3 years, 7 months ago
https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.S3.html
 Using Amazon S3 as a target for AWS Database Migration Service
 A
 upvoted 11 times

🗳️ 👤 **gelsm** 3 years, 7 months ago
https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.S3.html
 upvoted 1 times

🗳️ 👤 **novice_expert** Highly Voted 👍 3 years ago
Selected Answer: A
 -C and D out because of VPN Tunnel
 -B is out for multiple snowballs

A. 80 TB = 1 Snowball -> S3 -> DMS -> CDC -> Target
 upvoted 5 times

🗳️ 👤 **Pranava_GCP** Most Recent 🕒 1 year, 8 months ago
Selected Answer: A
 A. 80 TB = 1 Snowball -> S3 -> DMS -> CDC -> Target
 upvoted 2 times

🗳️ 👤 **Arun32** 2 years, 6 months ago
 A for sure. 2nd snowball device not required for CDC.
 upvoted 1 times

🗳️ 👤 **kret** 3 years, 1 month ago
Selected Answer: A
 80TB -> Snowball, ongoing changes rep with DMS
 upvoted 3 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago
Selected Answer: A
 A is the correct Answer

upvoted 4 times

🗨️ 👤 **thelad** 3 years, 3 months ago

Selected Answer: B

I'm taking B here as the company network bandwidth is pretty much redundant for use in the migration.

upvoted 3 times

🗨️ 👤 **Dip11** 3 years, 7 months ago

Ans: A

upvoted 2 times

🗨️ 👤 **Aesthet** 3 years, 7 months ago

A final answer

upvoted 3 times

🗨️ 👤 **manan728** 3 years, 7 months ago

You need snowball edge here. Of A and B only A makes more sense.

upvoted 2 times

🗨️ 👤 **wzlinux** 3 years, 7 months ago

A is for me.

upvoted 1 times

🗨️ 👤 **Xgab** 3 years, 8 months ago

My answer is A , D is incorrect SCT is for schemas and doesn't support CDC

upvoted 2 times

A database specialist is responsible for an Amazon RDS for MySQL DB instance with one read replica. The DB instance and the read replica are assigned to the default parameter group. The database team currently runs test queries against a read replica. The database team wants to create additional tables in the read replica that will only be accessible from the read replica to benefit the tests.

Which should the database specialist do to allow the database team to create the test tables?

- A. Contact AWS Support to disable read-only mode on the read replica. Reboot the read replica. Connect to the read replica and create the tables.
- B. Change the `read_only` parameter to false (`read_only=0`) in the default parameter group of the read replica. Perform a reboot without failover. Connect to the read replica and create the tables using the `local_only MySQL` option.
- C. Change the `read_only` parameter to false (`read_only=0`) in the default parameter group. Reboot the read replica. Connect to the read replica and create the tables.
- D. Create a new DB parameter group. Change the `read_only` parameter to false (`read_only=0`). Associate the read replica with the new group. Reboot the read replica. Connect to the read replica and create the tables.

Suggested Answer: A

Community vote distribution

D (100%)

🗳️ 👤 **jove** Highly Voted 3 years, 5 months ago

Selected Answer: D

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-read-replica/>
upvoted 7 times

🗳️ 👤 **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: D

D. Create new DB parameter group --> change parameter value --> associate replica with new param group --> reboot replica --> connect to read replica then create tables.
upvoted 2 times

🗳️ 👤 **navkumin** 1 year, 11 months ago

Selected Answer: D

Option D, as default parameter group cannot be modified.
upvoted 1 times

🗳️ 👤 **sk1974** 2 years, 3 months ago

Any reason, it's not B. <https://aws.amazon.com/premiumsupport/knowledge-center/rds-read-replica/>. The replica has to be rebooted without a failover option. However, I do not understand what it means 'local_only MySQL'.
upvoted 1 times

🗳️ 👤 **SeemaDataReader** 2 years ago

won't be B because default groups cannot be modified
upvoted 3 times

🗳️ 👤 **Dantas** 2 years, 11 months ago

Selected Answer: D

Option D
upvoted 2 times

🗳️ 👤 **KaranGandhi30** 3 years ago

Selected Answer: D

Option D
upvoted 2 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: D

`read_only` in default parameter group can not be changed (possible in custom parameter group)
<https://aws.amazon.com/premiumsupport/knowledge-center/rds-read-replica/>

upvoted 4 times

🗨️ 👤 **GMartinelli** 3 years, 5 months ago

Selected Answer: D

Option D

upvoted 2 times

🗨️ 👤 **guru_ji** 3 years, 7 months ago

This Q was asked in my exam.

upvoted 1 times

🗨️ 👤 **Hits_23** 3 years, 7 months ago

D is correct answer due to default parameter changes are not allowed

upvoted 1 times

🗨️ 👤 **Dip11** 3 years, 7 months ago

D is the correct ans.

upvoted 1 times

🗨️ 👤 **Aesthet** 3 years, 7 months ago

D final answer

upvoted 1 times

🗨️ 👤 **shyamnsingh** 3 years, 7 months ago

D For me

upvoted 3 times

🗨️ 👤 **Billhardy** 3 years, 7 months ago

D for me

upvoted 1 times

🗨️ 👤 **wzlinux** 3 years, 7 months ago

D is for me

upvoted 1 times

🗨️ 👤 **shantest1** 3 years, 8 months ago

I think D?

Default parameter group cannot be changed.

upvoted 3 times

🗨️ 👤 **shantest1** 3 years, 8 months ago

Sorry B.

upvoted 1 times

🗨️ 👤 **shantest1** 3 years, 7 months ago

Sorry C: - Do not think need to adjust anything like what is said in B:

* We can edit default parameter group (no need of extra step of creating additional parameter group)

upvoted 1 times

🗨️ 👤 **shantest1** 3 years, 7 months ago

Sorry again, default parameter cannot be changed.

D: answer.

No way to delete the comments, otherwise it will make it easy to update.

upvoted 3 times

A company has a heterogeneous six-node production Amazon Aurora DB cluster that handles online transaction processing (OLTP) for the core business and OLAP reports for the human resources department. To match compute resources to the use case, the company has decided to have the reporting workload for the human resources department be directed to two small nodes in the Aurora DB cluster, while every other workload goes to four large nodes in the same DB cluster.

Which option would ensure that the correct nodes are always available for the appropriate workload while meeting these requirements?

- A. Use the writer endpoint for OLTP and the reader endpoint for the OLAP reporting workload.
- B. Use automatic scaling for the Aurora Replica to have the appropriate number of replicas for the desired workload.
- C. Create additional readers to cater to the different scenarios.
- D. Use custom endpoints to satisfy the different workloads.

Suggested Answer: B

Community vote distribution

D (100%)

🗳️ 👤 **shantest1** Highly Voted 3 years, 7 months ago

D - Custom endpoints
upvoted 8 times

🗳️ 👤 **MultiAZ** Most Recent 1 year, 4 months ago

Selected Answer: D
D - custom endpoints can provide this flexibility
upvoted 1 times

🗳️ 👤 **examineme** 2 years, 5 months ago

Selected Answer: D
Use custom endpoints
upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: D
D. Use custom endpoints to satisfy the different workloads.

<https://aws.amazon.com/about-aws/whats-new/2018/11/amazon-aurora-simplifies-workload-management-with-custom-endpoints/>

you can now create custom endpoints for Amazon Aurora databases. This allows you to distribute and load balance workloads across different sets of database instances in your Aurora cluster.

upvoted 3 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

What does " heterogeneous six-node production Amazon Aurora DB cluster " even mean? I've never heard of any Aurora infrastructure described this way!!
upvoted 1 times

🗳️ 👤 **khchan123** 3 years ago

heterogeneous means the six nodes are of different size.
upvoted 3 times

🗳️ 👤 **guru_ji** 3 years, 7 months ago

Correct Answer: D
upvoted 1 times



🗳️ 👤 **damaldon** 3 years, 7 months ago

D is correct as per link:
<https://aws.amazon.com/about-aws/whats-new/2018/11/amazon-aurora-simplifies-workload-management-with-custom-endpoints/>
ou can now create custom endpoints for Amazon Aurora databases. This allows you to distribute and load balance workloads across different sets of

database instances in your Aurora cluster.



For example, you may provision a set of Aurora Replicas to use an instance type with higher memory capacity in order to run an analytics workload. A custom endpoint can then help you route the analytics workload to these appropriately-configured instances, while keeping other instances in your cluster isolated from this workload. As you add or remove instances from the custom endpoint to match your workload, the endpoint helps spread the load around.

upvoted 1 times

  **Hits_23** 3 years, 7 months ago

D is right choice. Custom endpoints. OLTP can only work with master instance endpoint. OLAP reports could work with either master instance/any combination of reader endpoints

upvoted 1 times

  **Aesthet** 3 years, 7 months ago

D final answer

upvoted 3 times

  **std2021** 3 years, 8 months ago

Answer D

upvoted 3 times

Developers have requested a new Amazon Redshift cluster so they can load new third-party marketing data. The new cluster is ready and the user credentials are given to the developers. The developers indicate that their copy jobs fail with the following error message:

`Amazon Invalid operation: S3ServiceException:Access Denied,Status 403,Error AccessDenied.`

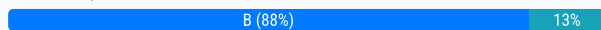
The developers need to load this data soon, so a database specialist must act quickly to solve this issue.

What is the MOST secure solution?

- A. Create a new IAM role with the same user name as the Amazon Redshift developer user ID. Provide the IAM role with read-only access to Amazon S3 with the assume role action.
- B. Create a new IAM role with read-only access to the Amazon S3 bucket and include the assume role action. Modify the Amazon Redshift cluster to add the IAM role.
- C. Create a new IAM role with read-only access to the Amazon S3 bucket with the assume role action. Add this role to the developer IAM user ID used for the copy job that ended with an error message.
- D. Create a new IAM user with access keys and a new role with read-only access to the Amazon S3 bucket. Add this role to the Amazon Redshift cluster. Change the copy job to use the access keys created.

Suggested Answer: D

Community vote distribution



std2021 Highly Voted 3 years, 8 months ago

Answer B

upvoted 6 times

jitesh_k Most Recent 1 year, 5 months ago

Why is C not correct?

upvoted 1 times

Hisayuki 1 year, 4 months ago

Because the COPY from S3 will be done by Redshift Cluster

upvoted 1 times

paws_brasil 1 year, 6 months ago

Selected Answer: D

Why it is not D ? It seems to be a question about Unload-Copy operation. The COPY operation requires access keys and secret keys.

upvoted 1 times

novice_expert 3 years ago

Selected Answer: B

B. new IAM role with read-only access to the Amazon S3 bucket and include the assume role action.

Modify the Amazon Redshift cluster to add the IAM role.

Provide authentication for your cluster to access Amazon S3 on your behalf to load the sample data. You provide authentication by referencing the IAM role that you created and set as the default for your cluster

upvoted 4 times

user0001 3 years, 3 months ago

B

<https://docs.aws.amazon.com/redshift/latest/gsg/rs-gsg-create-sample-db.html>

Provide authentication for your cluster to access Amazon S3 on your behalf to load the sample data. You provide authentication by referencing the IAM role that you created and set as the default for your cluster

upvoted 3 times

GMartinelli 3 years, 5 months ago

Selected Answer: B

Option B

upvoted 3 times

🗨️ 👤 **ChauPhan** 3 years, 6 months ago

B. To best protect your sensitive data and safeguard your AWS access credentials, we recommend creating an IAM role and attaching it to your cluster. For more information about providing access permissions, see [Permissions to access other AWS resources](#).

In this step, you create a new IAM role that allows Amazon Redshift to load data from Amazon S3 buckets. An IAM role is an IAM identity that you can create in your account that has specific permissions. In the next step, you attach the role to your cluster

upvoted 1 times

🗨️ 👤 **guru_ji** 3 years, 7 months ago

anyone here already cleared the exam ?

How much %age Q we will get in real exam, any idea ?

upvoted 2 times

🗨️ 👤 **gelsm** 3 years, 7 months ago

Answer: B

<https://docs.aws.amazon.com/redshift/latest/gsg/rs-gsg-create-an-iam-role.html>

"Now that you have created the new role, your next step is to attach it to your cluster. You can attach the role when you launch a new cluster or you can attach it to an existing cluster. In the next step, you attach the role to a new cluster."

upvoted 2 times

🗨️ 👤 **Hits_23** 3 years, 7 months ago

Correct answer is B. Good practice is to create an IAM Role with read only permission to S3 and attach this role to Redshift cluster for COPY jobs from S3 to Redshift cluster

upvoted 2 times

🗨️ 👤 **Aesthet** 3 years, 7 months ago

B

<https://docs.aws.amazon.com/redshift/latest/gsg/rs-gsg-create-an-iam-role.html>

upvoted 1 times

🗨️ 👤 **Aesthet** 3 years, 7 months ago

Role-based access control

https://docs.aws.amazon.com/redshift/latest/dg/copy-usage_notes-access-permissions.html

upvoted 3 times

🗨️ 👤 **manan728** 3 years, 7 months ago

C is the correct one. You cannot modify the redshift cluster to add IAM role to it. But you can add the role to the user who's trying to get access to S3

upvoted 1 times

🗨️ 👤 **[Removed]** 3 years, 7 months ago

you can modify the redshift cluster to associal an IAM role. C is not correct

upvoted 5 times

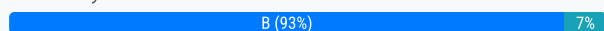
A database specialist at a large multi-national financial company is in charge of designing the disaster recovery strategy for a highly available application that is in development. The application uses an Amazon DynamoDB table as its data store. The application requires a recovery time objective (RTO) of 1 minute and a recovery point objective (RPO) of 2 minutes.

Which operationally efficient disaster recovery strategy should the database specialist recommend for the DynamoDB table?

- A. Create a DynamoDB stream that is processed by an AWS Lambda function that copies the data to a DynamoDB table in another Region.
- B. Use a DynamoDB global table replica in another Region. Enable point-in-time recovery for both tables.
- C. Use a DynamoDB Accelerator table in another Region. Enable point-in-time recovery for the table.
- D. Create an AWS Backup plan and assign the DynamoDB table as a resource.

Suggested Answer: C

Community vote distribution



shantest1 Highly Voted 3 years, 8 months ago

B DynamoDB global tables

upvoted 16 times

Suresh108 Highly Voted 3 years, 6 months ago

BBBBBBBBBBB

global tables are multi master replication enabled solution.

it meets the requirement.

upvoted 6 times

IhorK Most Recent 1 year, 9 months ago

Selected Answer: B

- A single Amazon DynamoDB global table can only have one replica table per AWS Region.

- You can enable point-in-time recovery on each replica of a global table.

https://aws.amazon.com/dynamodb/global-tables/?nc1=h_ls

upvoted 2 times

lollyj 2 years, 5 months ago

Selected Answer: B

Only logical answer. Backups take too long to restore. Global tables are synchronously updated

upvoted 1 times

sachin 2 years, 10 months ago

It is asking for DR solution. My vote is for A . If you need less RPO and RTO setup Streams and Kinesis, Lambda and you can achieve the given RPO RTO,

RPO is 5 mins in usual automated backups . B could be true but in global tables also RPO is 5 mins

upvoted 2 times

Jiang_aws1 2 years, 7 months ago

B is correct. Global table is multi master & it is sync up with few sec

upvoted 2 times

Dantas 2 years, 12 months ago

Selected Answer: B

Two-minute recovery point objective (RPO).

upvoted 2 times

novice_expert 3 years ago

Selected Answer: B

x A. Create a DynamoDB stream that is processed by an AWS Lambda function that copies the data to a DynamoDB table in another Region. (good if B is wrong)

B. Use a DynamoDB global table replica in another Region. Enable point-in-time recovery for both tables. (PITR is distraction btw PITR takes 5 min

recovery, but global table will auto failover to good region, right?)

x C. DynamoDB Accelerator (its for read caching)

x D. Create an AWS Backup plan and assign the DynamoDB table as a resource. (takes 1 hour)

upvoted 4 times

🗨️ **marcoeu** 3 years ago

Selected Answer: B

B because of RTO and RPO limits defined. D is good solution doesn't fit limits. DynamoDB Accelerator is distractor for DR question.

upvoted 2 times

🗨️ **Sandy1971** 3 years, 4 months ago

Selected Answer: D

A, B and C are not DR solutions for Dynamo DB. Global Tables is multi region replication.

D meets the objective . You can set up scheduled backups for Amazon DynamoDB using AWS Backup.. <https://aws.amazon.com/blogs/database/set-up-scheduled-backups-for-amazon-dynamodb-using-aws-backup/>

upvoted 1 times

🗨️ **RotterDam** 3 years, 2 months ago

to get RPO of 2 minutes you will have to create a backup plan that backs up < 2 minutes!! Thats not what backup is for

upvoted 5 times

🗨️ **Shunpin** 3 years, 5 months ago

Selected Answer: B

I will go to B.

<https://s3.amazonaws.com/solutions-reference/multi-region-application-architecture/latest/multi-region-application-architecture.pdf>

upvoted 2 times

🗨️ **toppic26** 3 years, 6 months ago

A is the answer.

D is wrong. Restore takes nearly 1 hour https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/backuprestore_HowItWorks.html

B is wrong. RPO is about 5 mins.

A is true. If you need less RPO and RTO setup Streams and Kinesis, Lambda whatever

upvoted 4 times

🗨️ **Sp230** 3 years, 6 months ago

The question requires 2 min RPO. Using PITR RPO is about 5 min. I think D makes more sense here

upvoted 1 times

🗨️ **toppic26** 3 years, 6 months ago

But restore is "less than one hour" aws says

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

upvoted 1 times

🗨️ **Scunningham99** 3 years, 6 months ago

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

upvoted 2 times

🗨️ **Aesthet** 3 years, 7 months ago

B final answer

upvoted 4 times

🗨️ **gelsm** 3 years, 7 months ago

Hi Aesthet, How are you sure about this? :) Thanks!

upvoted 1 times

A small startup company is looking to migrate a 4 TB on-premises MySQL database to AWS using an Amazon RDS for MySQL DB instance. Which strategy would allow for a successful migration with the LEAST amount of downtime?

- A. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data center. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket. Import the snapshot into the DB instance utilizing the MySQL utilities running on an Amazon EC2 instance. Immediately point the application to the DB instance.
- B. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data center. Use the mysqldump utility to create a snapshot of the on-premises MySQL server. Copy the snapshot into the EC2 instance and restore it into the EC2 MySQL instance. Use AWS DMS to migrate data into a new RDS for MySQL DB instance. Point the application to the DB instance.
- C. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data center. Use the mysqldump utility to create a snapshot of the on-premises MySQL server. Copy the snapshot into an Amazon S3 bucket and import the snapshot into a new RDS for MySQL DB instance using the MySQL utilities running on an EC2 instance. Point the application to the DB instance.
- D. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data center. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucket. Import the snapshot into the DB instance using the MySQL utilities running on an Amazon EC2 instance. Establish replication into the new DB instance using MySQL replication. Stop application access to the on-premises MySQL server and let the remaining transactions replicate over. Point the application to the DB instance.

Suggested Answer: B

Community vote distribution

D (88%)

13%

 **novice_expert** Highly Voted 3 years ago

Selected Answer: D

on-premises MySQL server -> mysqldump utility -> snapshot -> copy to S3 -> MySQL Utility on EC2 -> import to AWS RDS MySQL -> Establish replication into the new DB instance using MySQL replication -> Stop application access to the on-premises MySQL server -> let the remaining transactions replicate over. -> Point the application to the DB instance.

upvoted 5 times

 **aws2023a** Most Recent 1 year, 9 months ago

It says RDS so B, C (EC2 option) is out.

upvoted 1 times

 **megramlak** 1 year, 11 months ago

I feel correct answer is B as DMS is involved


upvoted 1 times

 **lollyj** 2 years, 5 months ago

Selected Answer: D

Was a toss between B&D. DMS task threw me off but I don't see how you can get snapshot to an EC2 instance from option B

upvoted 2 times

 **vkruiger** 2 years, 5 months ago

Selected Answer: B

Correct answer is B

I don't get why so many people chose D. How is extra copy to S3 bucket reduce downtime?

upvoted 1 times

 **lollyj** 2 years, 5 months ago

Installing new SQL software on an EC2 just for the sake of migration threw me off. Getting that much data through an s3 gateway endpoint would get the data from on-prem to the VPC but directly on to an EC2 instance seems ridic to me. So answer is D



upvoted 1 times

 **damaldon** 3 years, 7 months ago

D looks correct as per this link:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/MySQL.Procedural.Importing.NonRDSRepl.html>

upvoted 3 times

  **Suresh108** 3 years, 7 months ago

DDDDDDDD

upvoted 1 times

  **Suresh108** 3 years, 7 months ago

Ignore all EC2 based answers - method of elimination - left with D

upvoted 1 times

  **Aesthet** 3 years, 7 months ago

D

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/MySQL.Procedural.Importing.External.Repl.html>

upvoted 2 times

  **shantest1** 3 years, 7 months ago

D: Ans

upvoted 2 times

  **std2021** 3 years, 8 months ago

Option D

upvoted 2 times

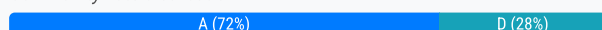
A software development company is using Amazon Aurora MySQL DB clusters for several use cases, including development and reporting. These use cases place unpredictable and varying demands on the Aurora DB clusters, and can cause momentary spikes in latency. System users run ad-hoc queries sporadically throughout the week. Cost is a primary concern for the company, and a solution that does not require significant rework is needed.

Which solution meets these requirements?

- A. Create new Aurora Serverless DB clusters for development and reporting, then migrate to these new DB clusters.
- B. Upgrade one of the DB clusters to a larger size, and consolidate development and reporting activities on this larger DB cluster.
- C. Use existing DB clusters and stop/start the databases on a routine basis using scheduling tools.
- D. Change the DB clusters to the burstable instance family.

Suggested Answer: D

Community vote distribution



🗳️ 👤 **Zhongkai** Highly Voted 3 years, 7 months ago

Check <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Concepts.DBInstanceClass.html> - t2 and t3 (8vCPU 32GB RAM) series are mostly weaker than r4/r5 (96 vCPU 768GB RAM) series. Changing from r to t series is not a good idea. Furthermore, "unpredictable" indicates "Aurora Serverless". Hence I will go with A.

upvoted 15 times

🗳️ 👤 **novice_expert** Highly Voted 3 years ago

Selected Answer: A

"unexpected and variable demands, adhoc searched" => "Aurora Serverless"

upvoted 7 times

🗳️ 👤 **Germaneli** Most Recent 1 year, 8 months ago

Selected Answer: A

Skipping B+c as distractors.

Torn between A+D -

D is not applicable as "we recommend using the [burstable] T DB instance classes only for development, test, or other nonproduction servers".

Switching to burstable classes is not an option here for PROD reporting.

Hence A.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.BestPractices.html#AuroraMySQL.BestPractices.T2Medium>

upvoted 1 times

🗳️ 👤 **anantarb** 2 years, 3 months ago

Selected Answer: D

Question says without significant rework. Option A would require migration. So option D is correct.

<https://aws.amazon.com/rds/instance-types/>

upvoted 4 times

🗳️ 👤 **lollyj** 2 years, 5 months ago

Selected Answer: D

A - Serverless will require moving existing DB although you are saving on cost.

D. Can respond to traffic spikes, requires very little change and the costs for the spikes won't be much because it isn't too frequent. The burst credits can be used.

upvoted 1 times

🗳️ 👤 **Dantas** 3 years, 2 months ago

Selected Answer: A

<https://aws.amazon.com/rds/aurora/serverless/>

Amazon Aurora Serverless is a simple, cost-effective option for infrequent, intermittent, or unpredictable workloads.

upvoted 4 times

🗄️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: A

ad-hoc searches = Serverless

<https://aws.amazon.com/rds/aurora/faqs/?nc=sn&loc=6>

Q: Can I migrate an existing Aurora DB cluster to Aurora Serverless?

Yes, you can restore a snapshot taken from an existing Aurora provisioned cluster into an Aurora Serverless DB Cluster (and vice versa).

upvoted 1 times

🗄️ 👤 **user0001** 3 years, 3 months ago

D

A is good if the cluster experience no activities which is not the case

upvoted 2 times

🗄️ 👤 **awsmonster** 3 years, 4 months ago

Answer: A

The company's key concern is cost, and a solution that does not need extensive rework is required. Converting Aurora Cluster to Serverless does not require extensive rework.

upvoted 1 times

🗄️ 👤 **SMAZ** 3 years, 5 months ago

"resulting in brief latency spikes" Answer is D

upvoted 2 times

🗄️ 👤 **jove** 3 years, 6 months ago

If the key concern is the cost, the right option is Aurora Serverless.. Answer is A

upvoted 2 times

🗄️ 👤 **Scunningham99** 3 years, 6 months ago

D less rework - A is too much work

upvoted 5 times

🗄️ 👤 **guru_ji** 3 years, 7 months ago

Answer: D

****does not require significant rework is needed****

upvoted 2 times

🗄️ 👤 **CW0106** 3 years, 7 months ago

D. no significant work.

upvoted 2 times

🗄️ 👤 **ChauPhan** 3 years, 7 months ago

A. Aurora Serverless is most saving cost. You pay what you query

upvoted 1 times

🗄️ 👤 **Hits_23** 3 years, 7 months ago

Answer is A. Serverless aurora cluster can be cost effective in the varying workload and dev/test environment where services are not required 24 * 7

upvoted 1 times

🗄️ 👤 **Aesthet** 3 years, 7 months ago

A final answer

upvoted 3 times

A database specialist is building a system that uses a static vendor dataset of postal codes and related territory information that is less than 1 GB in size. The dataset is loaded into the application's cache at start up. The company needs to store this data in a way that provides the lowest cost with a low application startup time.

Which approach will meet these requirements?

- A. Use an Amazon RDS DB instance. Shut down the instance once the data has been read.
- B. Use Amazon Aurora Serverless. Allow the service to spin resources up and down, as needed.
- C. Use Amazon DynamoDB in on-demand capacity mode.
- D. Use Amazon S3 and load the data from flat files.

Suggested Answer: A

Community vote distribution

D (67%)

C (33%)

🗳️ 👤 **TonyGe** Highly Voted 3 years, 7 months ago

D, key words " static vendor dataset" & "lowest cost"
upvoted 7 times

🗳️ 👤 **MultiAZ** Most Recent 1 year, 4 months ago

Selected Answer: D

D gives the lower cost. It also meets the requirements - S3 can do tons of IOPS when needed. Most importantly, the data is loaded on startup and then used by the app in-memory, not used constantly - so there is no need for DynamoDB.
upvoted 1 times

🗳️ 👤 **IhorK** 1 year, 9 months ago

Selected Answer: C

DynamoDB in on-demand capacity mode price:

- DynamoDB Standard table class, Read Request Units (RRU) - \$0.25 per MILLION read request units.
- Data storage - First 25 GB stored per month is free using the DynamoDB Standard table class.
- On-demand backup - Cold Backup Storage* \$0.03 per GB-month (The information is static, you can do without a backup).

<https://aws.amazon.com/dynamodb/pricing/on-demand/>

Amazon S3 pricing:

S3 Standard - General purpose storage for any type of data, typically used for frequently accessed data. First 50 TB / Month - \$0.023 per GB

<https://aws.amazon.com/s3/pricing/>

Application does not start so often, I would choose answer C. DynamoDB will give this information much faster than reading from S3.

The price of \$0.25 per MILLION read request is cheaper than S3.

upvoted 2 times

🗳️ 👤 **[Removed]** 1 year, 11 months ago

Selected Answer: C

the lowest cost & a low application startup time -> Dynamo DB with OnDemand capacity mode.

Option D (S3) will cause longer startup time in my opinion

upvoted 2 times

🗳️ 👤 **satishstechie** 2 years, 5 months ago

Selected Answer: D

simple option

upvoted 1 times

🗳️ 👤 **lollyj** 2 years, 5 months ago

Selected Answer: D

Static small files. S3 is your best bet for costs

upvoted 2 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: D

x A. RDS (costly for 1 GB data)
x B. Amazon Aurora Serverless. (costly for 1 GB data)
x C. Use Amazon DynamoDB in on-demand capacity mode.
(postal codes and associated territorial data that is less than 1 GB in size => key value data,
minimizing application launch time => dynamoDB, but caching is required at application and not database)

D. Use Amazon S3 and load the data from flat files. (most cost-effective)
upvoted 4 times

🗳️ 👤 **Scunningham99** 3 years, 6 months ago
D is the answer
upvoted 2 times

🗳️ 👤 **aws4myself** 3 years, 6 months ago
Let us see in this way, 1 GB file on boot uploading in to application time < (1 GB file from s3 to Application + load time). So definitely it is not D, I will go with B
upvoted 1 times

🗳️ 👤 **gelsm** 3 years, 7 months ago
I agree with D based on the link below:
<https://www.sumologic.com/insight/s3-cost-optimization/>
For example, for 1 GB file stored on S3 with 1 TB of storage provisioned, you are billed for 1 GB only. In a lot of other services such as Amazon EC2, Amazon Elastic Block Storage (Amazon EBS) and Amazon DynamoDB you pay for provisioned capacity. For example, in the case of Amazon EBS disk you pay for the size of 1 TB of disk even if you just save 1 GB file. This makes managing S3 cost easier than many other services including Amazon EBS and Amazon EC2. On S3 there is no risk of over-provisioning and no need to manage disk utilization.
upvoted 1 times

🗳️ 👤 **Suresh108** 3 years, 7 months ago
DDDDDDDDDD. lowest cost.

--> this says Application cache and NOT DB cache.

application's cache at start up
upvoted 2 times

🗳️ 👤 **AM** 3 years, 7 months ago
OR D is also a possibility
upvoted 2 times

🗳️ 👤 **AM** 3 years, 7 months ago
I will go with B. Will cover lowest cost.
upvoted 1 times

🗳️ 👤 **Aesthet** 3 years, 7 months ago
C or D
I choose D
upvoted 2 times

🗳️ 👤 **shantest1** 3 years, 7 months ago
C: ?
lowest start up time asked
upvoted 1 times

🗳️ 👤 **shantest1** 3 years, 7 months ago
May be D:
Application Cache may be enough. Load from S3 straight into Application Cache
upvoted 6 times

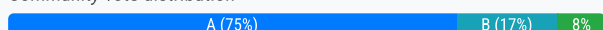
A database specialist needs to review and optimize an Amazon DynamoDB table that is experiencing performance issues. A thorough investigation by the database specialist reveals that the partition key is causing hot partitions, so a new partition key is created. The database specialist must effectively apply this new partition key to all existing and new data.

How can this solution be implemented?

- A. Use Amazon EMR to export the data from the current DynamoDB table to Amazon S3. Then use Amazon EMR again to import the data from Amazon S3 into a new DynamoDB table with the new partition key.
- B. Use AWS DMS to copy the data from the current DynamoDB table to Amazon S3. Then import the DynamoDB table to create a new DynamoDB table with the new partition key.
- C. Use the AWS CLI to update the DynamoDB table and modify the partition key.
- D. Use the AWS CLI to back up the DynamoDB table. Then use the `restore-table-from-backup` command and modify the partition key.

Suggested Answer: D

Community vote distribution



shantest1 Highly Voted 3 years, 7 months ago

A. Answer

Use EMR to copy to S3 and use EMR to create new table.

upvoted 9 times

Suresh108 Highly Voted 3 years, 7 months ago

AAAA for me as well.

<https://aws.amazon.com/premiumsupport/knowledge-center/back-up-dynamodb-s3/>

EMR to backup and restore

upvoted 5 times

aviathor Most Recent 1 year, 11 months ago

Selected Answer: A

A. Use Amazon EMR to export the data from the current DynamoDB table to Amazon S3. Then use Amazon EMR again to import the data from Amazon S3 into a new DynamoDB table with the new partition key. Most Voted

xB. DynamoDB is not a valid source for AWS DMS

xC. The partition key cannot be modified

xD. Although you can choose to restore a DynamoDB table without the secondary indexes, the partition key cannot be modifier during restore

upvoted 3 times

ninjalight25 2 years, 2 months ago

Selected Answer: C

To apply the new partition key to all existing and new data in a DynamoDB table, the AWS CLI can be used to update the table and modify the partition key

upvoted 1 times

lollyj 2 years, 5 months ago

Selected Answer: B

I believe DMS works with DynamoDB and is a simpler solution compared to EMR. Dynamodb can be a source for DMS.

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.DynamoDB.html

upvoted 2 times

Kanwar_89 2 years, 4 months ago

Dynamodb can be a target , but not a source -- Hence A

upvoted 1 times

Chirantan 2 years, 10 months ago

Is it possible to change partition key in DynamoDB? No. Once the table is setup, you cannot modify its Key Schema. You can only provision a new table, move data there, and then remove the first table

upvoted 5 times

🗨️ **novice_expert** 3 years ago

Selected Answer: A

<https://aws.amazon.com/premiumsupport/knowledge-center/back-up-dynamodb-s3/>

DynamoDB to S3 is possible by EMR, AWS Glue, or AWS Data Pipeline.

upvoted 3 times

🗨️ **tugboat** 3 years, 3 months ago

Selected Answer: A

DynamoDB cannot be a source for DMS.

Per - <https://dynobase.com/dynamodb-keys/>

Is it possible to change partition key in DynamoDB?

No. Once the table is setup, you cannot modify its Key Schema. You can only provision a new table, move data there, and then remove the first table.

Per - <https://aws.amazon.com/premiumsupport/knowledge-center/back-up-dynamodb-s3/>

To customize the process of creating backups, you can use use Amazon EMR, AWS Glue, or AWS Data Pipeline.

So, using EMR is the only valid option.

upvoted 3 times

🗨️ **kped21** 3 years, 3 months ago

D

<https://docs.aws.amazon.com/cli/latest/reference/dynamodb/restore-table-from-backup.html>

upvoted 1 times

🗨️ **RotterDam** 3 years, 2 months ago

this will recreate the same partition key

upvoted 1 times

🗨️ **Aesthet** 3 years, 7 months ago

A, other options are not possible

upvoted 3 times

🗨️ **std2021** 3 years, 8 months ago

I'll go with B.

restore-table-from-backup doesn't seem to be able to change partition key, not sure here

upvoted 2 times

🗨️ **toppic26** 3 years, 6 months ago

DynamoDB doesnt seem to be a source for DMS. https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Source.html

upvoted 4 times

A company is going through a security audit. The audit team has identified cleartext master user password in the AWS CloudFormation templates for Amazon

RDS for MySQL DB instances. The audit team has flagged this as a security risk to the database team.

What should a database specialist do to mitigate this risk?

- A. Change all the databases to use AWS IAM for authentication and remove all the cleartext passwords in CloudFormation templates.
- B. Use an AWS Secrets Manager resource to generate a random password and reference the secret in the CloudFormation template.
- C. Remove the passwords from the CloudFormation templates so Amazon RDS prompts for the password when the database is being created.
- D. Remove the passwords from the CloudFormation template and store them in a separate file. Replace the passwords by running CloudFormation using a sed command.

Suggested Answer: C

Community vote distribution

B (100%)

🗳️ 👤 **ImprovMAN** 2 years, 6 months ago

Answer B

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: B

Use an AWS Secrets Manager resource to generate a random password and reference the secret in the CloudFormation template.

upvoted 4 times

🗳️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: B

secret manager to the rescue!

upvoted 3 times

🗳️ 👤 **soyyodario** 3 years, 3 months ago

Selected Answer: B

Answer B

upvoted 2 times

🗳️ 👤 **GMartinelli** 3 years, 5 months ago

Selected Answer: B

Option B

upvoted 2 times

🗳️ 👤 **ChauPhan** 3 years, 6 months ago

B no doubt

upvoted 1 times

🗳️ 👤 **gelsm** 3 years, 6 months ago

Answer: B

<https://aws.amazon.com/blogs/infrastructure-and-automation/securing-passwords-in-aws-quick-starts-using-aws-secrets-manager/>

Saving a password in a clear text file is not a secure practice.

Today, I want to discuss how you can store secrets in Secrets Manager via AWS CloudFormation. Then I'll show, using code examples, how to retrieve secrets.

upvoted 3 times



🗳️ 👤 **AM** 3 years, 7 months ago



Even though B works, A will also work as it will generate token for login. Also there is no info if rotation of password is required. I am not sure between A & B.



upvoted 2 times

🗳️ 👤 **rlnd2000** 2 years, 10 months ago

It is not necessary to state in the question that the rotation of passwords for services accounts, such as master user, it is a best practice.
upvoted 1 times

  **Aesthet** 3 years, 7 months ago
B final answer
upvoted 2 times

  **shantest1** 3 years, 8 months ago
B. Ans
upvoted 4 times

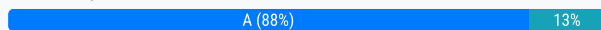
  **std2021** 3 years, 8 months ago
Option B
upvoted 2 times

A company's database specialist disabled TLS on an Amazon DocumentDB cluster to perform benchmarking tests. A few days after this change was implemented, a database specialist trainee accidentally deleted multiple tables. The database specialist restored the database from available snapshots. An hour after restoring the cluster, the database specialist is still unable to connect to the new cluster endpoint. What should the database specialist do to connect to the new, restored Amazon DocumentDB cluster?

- A. Change the restored cluster's parameter group to the original cluster's custom parameter group.
- B. Change the restored cluster's parameter group to the Amazon DocumentDB default parameter group.
- C. Configure the interface VPC endpoint and associate the new Amazon DocumentDB cluster.
- D. Run the syncInstances command in AWS DataSync.

Suggested Answer: B

Community vote distribution



std2021 Highly Voted 3 years, 8 months ago

Option A

upvoted 9 times

Pranava_GCP Most Recent 1 year, 8 months ago

Selected Answer: A

A. Change the restored cluster's parameter group to the original cluster's custom parameter group.

upvoted 1 times

Pranava_GCP 1 year, 8 months ago

<https://docs.aws.amazon.com/documentdb/latest/developerguide/security.encryption.ssl.html>

upvoted 1 times

megramlak 1 year, 11 months ago

A looks good, as TLS was disabled

upvoted 1 times

Paulv82003 1 year, 11 months ago

There is no Amazon DocumentDB VPC endpoint, open console and check. Amazon DocumentDB belongs in a VPC and the service endpoints are created with it.

upvoted 3 times

Kamalt 2 years, 1 month ago

ChatGPT:

If the restored Amazon DocumentDB cluster endpoint is not accessible, the database specialist should check the cluster's security group settings to ensure that the appropriate inbound rules are configured to allow incoming connections.

Therefore, none of the options presented in the question would solve the connectivity issue.

A and B refer to the cluster's parameter group, which does not affect connectivity to the cluster endpoint. Parameter groups are used to configure database engine settings.

upvoted 2 times

Kamalt 2 years, 1 month ago

C suggests configuring the interface VPC endpoint and associating the new Amazon DocumentDB cluster. However, this option is not relevant to the issue at hand as it relates to network connectivity within the VPC.

D suggests using AWS DataSync to run the syncInstances command, which is also not relevant to this issue. DataSync is a data transfer service, and running the syncInstances command does not solve connectivity issues to an Amazon DocumentDB cluster.

Therefore, the database specialist should verify the security group rules for the Amazon DocumentDB cluster to ensure that the necessary inbound rules are configured to allow incoming connections

upvoted 1 times

🗨️ 👤 **tucobbad** 2 years, 5 months ago

Selected Answer: A

Option A because context is TLS was disabled to perform benchmark testing so a new parameter group different from the default was created. As the default is to have TLS enabled, it will not be possible to connect to a new restored database with default settings so the original, previous custom parameter group should be assigned to it.

upvoted 4 times

🗨️ 👤 **lollyj** 2 years, 5 months ago

Selected Answer: C

Can someone explain why C isn't the answer? I understand why option A could be an answer but why not C?

upvoted 1 times

🗨️ 👤 **Paulv82003** 1 year, 11 months ago

There is no Amazon DocumentDB VPC endpoint, open console and check. Amazon DocumentDB belongs in a VPC and the service endpoints are created with it.

upvoted 1 times

🗨️ 👤 **novice_expert** 3 years ago

Selected Answer: A

A. Change the restored cluster's parameter group to the original cluster's custom parameter group.

Trick question: You can't modify the parameter settings of the default parameter groups.

You can use a DB parameter group to act as a container for engine configuration values that are applied to one or more DB instances.

upvoted 2 times

🗨️ 👤 **guru_ji** 3 years, 7 months ago

Correct Answer: A

upvoted 1 times

🗨️ 👤 **gelsm** 3 years, 7 months ago

A

Trick question: You can't modify the parameter settings of the default parameter groups.

You can use a DB parameter group to act as a container for engine configuration values that are applied to one or more DB instances.

If you create a DB instance without specifying a DB parameter group, the DB instance uses a default DB parameter group. Each default DB parameter group contains database engine defaults and Amazon RDS system defaults. You can't modify the parameter settings of a default parameter group. Instead, you create your own parameter group where you choose your own parameter settings. Not all DB engine parameters can be changed in a parameter group that you create.

upvoted 4 times

🗨️ 👤 **Suresh108** 3 years, 7 months ago

AAAA is looking good

upvoted 1 times

🗨️ 👤 **Aesthet** 3 years, 7 months ago

A final answer

upvoted 2 times

A company runs a customer relationship management (CRM) system that is hosted on-premises with a MySQL database as the backend. A custom stored procedure is used to send email notifications to another system when data is inserted into a table. The company has noticed that the performance of the CRM system has decreased due to database reporting applications used by various teams. The company requires an AWS solution that would reduce maintenance, improve performance, and accommodate the email notification feature. Which AWS solution meets these requirements?

- A. Use MySQL running on an Amazon EC2 instance with Auto Scaling to accommodate the reporting applications. Configure a stored procedure and an AWS Lambda function that uses Amazon SES to send email notifications to the other system.
- B. Use Amazon Aurora MySQL in a multi-master cluster to accommodate the reporting applications. Configure Amazon RDS event subscriptions to publish a message to an Amazon SNS topic and subscribe the other system's email address to the topic.
- C. Use MySQL running on an Amazon EC2 instance with a read replica to accommodate the reporting applications. Configure Amazon SES integration to send email notifications to the other system.
- D. Use Amazon Aurora MySQL with a read replica for the reporting applications. Configure a stored procedure and an AWS Lambda function to publish a message to an Amazon SNS topic. Subscribe the other system's email address to the topic.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **shantest1** Highly Voted 3 years, 8 months ago

D. Answer

Reduce maintenance, hosting database on EC2 increases maintenance.

upvoted 8 times

🗳️ 👤 **manan728** 3 years, 7 months ago

Did you take the real test @shantest1? Are the questions here valid if you know by any chance.

upvoted 1 times

🗳️ 👤 **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: D

D. Use Amazon Aurora MySQL with a read replica for the reporting applications. Configure a stored procedure and an AWS Lambda function to publish a message to an Amazon SNS topic. Subscribe the other system's email address to the topic.

upvoted 2 times

🗳️ 👤 **Pranava_GCP** 1 year, 8 months ago

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Integrating.Lambda.html>

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: D

reporting => read replica

upvoted 2 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

Got this question in my exam. (i cleared it). D is correct

upvoted 2 times

🗳️ 👤 **GMartinelli** 3 years, 5 months ago

Selected Answer: D

Option D

upvoted 1 times

🗳️ 👤 **guru_ji** 3 years, 6 months ago



Correct Answer: D

upvoted 1 times

🗳️ 👤 **Suresh108** 3 years, 6 months ago

DDDDDDDD is my choice

upvoted 1 times

  **Aesthet** 3 years, 7 months ago

D final answer

upvoted 2 times

  **Zhongkai** 3 years, 7 months ago

RDS event subscriptions do not cover "data is inserted into a table" - see

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_Events.Messages.html

We can use stored procedure to invoke Lambda function -

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Integrating.Lambda.html>

So I will go with B

upvoted 2 times

  **Zhongkai** 3 years, 7 months ago

Typo. I mean D

upvoted 1 times

A company needs to migrate Oracle Database Standard Edition running on an Amazon EC2 instance to an Amazon RDS for Oracle DB instance with Multi-AZ.

The database supports an ecommerce website that runs continuously. The company can only provide a maintenance window of up to 5 minutes. Which solution will meet these requirements?

- A. Configure Oracle Real Application Clusters (RAC) on the EC2 instance and the RDS DB instance. Update the connection string to point to the RAC cluster. Once the EC2 instance and RDS DB instance are in sync, fail over from Amazon EC2 to Amazon RDS.
- B. Export the Oracle database from the EC2 instance using Oracle Data Pump and perform an import into Amazon RDS. Stop the application for the entire process. When the import is complete, change the database connection string and then restart the application.
- C. Configure AWS DMS with the EC2 instance as the source and the RDS DB instance as the destination. Stop the application when the replication is in sync, change the database connection string, and then restart the application.
- D. Configure AWS DataSync with the EC2 instance as the source and the RDS DB instance as the destination. Stop the application when the replication is in sync, change the database connection string, and then restart the application.

Suggested Answer: B

Reference:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_UpgradeDBInstance.Oracle.html

Community vote distribution

C (100%)

manan728 **Highly Voted** 3 years, 7 months ago

C is most plausible option even though it does not mention change data capture or CDC. Rest of the steps seem legit in option C. No other given alternative would ideally meet the 5 min downtime requirement.

upvoted 13 times

RotterDam **Highly Voted** 3 years, 2 months ago

Selected Answer: C

this question came up in my exam. I chose (C). I passed the exam - also fairly certain this is the right answer. (I'll add this in as many questions as possible)

upvoted 10 times

Pranava_GCP **Most Recent** 1 year, 8 months ago

Selected Answer: C

C. Configure AWS DMS with the EC2 instance as the source and the RDS DB instance as the destination. Stop the application when the replication is in sync, change the database connection string, and then restart the application.

upvoted 1 times

IhorK 1 year, 9 months ago

Selected Answer: C

According to

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/migrate-an-on-premises-oracle-database-to-amazon-rds-for-oracle.html>

B and C could be the answer

- AWS Database Migration Service (AWS DMS);

- Native Oracle tools.

But B says "Stop the application for the entire process", while we have "a maintenance window of up to 5 minutes".

So, the answer is C.

upvoted 1 times

lollyj 2 years, 5 months ago

Selected Answer: C

C is the most plausible

upvoted 1 times

novice_expert 3 years ago

Selected Answer: C

x B. export import would take long time

C. EC2 -> DMS -> new RDS -> Stop when in sync.

Change db connections

upvoted 4 times

🗳️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: C

C as DMS enables fastest failover

upvoted 2 times

🗳️ 👤 **kped21** 3 years, 3 months ago

C: Correct

B - Invalid as the exp/imp can take long time and there is only 5 min downtime

upvoted 2 times

🗳️ 👤 **Shunpin** 3 years, 5 months ago

Ans:C

Oracle SE only support single thread export/import. Not possible to finish within 5 minutes.

upvoted 2 times

🗳️ 👤 **jove** 3 years, 6 months ago

C is correct

upvoted 2 times

🗳️ 👤 **toppic26** 3 years, 6 months ago

AWS Question with DMS, solution is Oracle tool, are you serious? Answer is C

upvoted 3 times

🗳️ 👤 **Scunningham99** 3 years, 6 months ago

CCCCC -- B is a big no even for a small database would take more than 5 mins. Stopping the application for the entire process etc

upvoted 2 times

🗳️ 👤 **guru_ji** 3 years, 6 months ago

Correct Answer: C

upvoted 1 times

🗳️ 👤 **Suresh108** 3 years, 7 months ago

only possible choices B and C.

B is wrong -- datapump import method takes longer time to migrate. it can't be done in 5 minutes - forget about it.

My choice is C.

upvoted 1 times

🗳️ 👤 **Aesthet** 3 years, 7 months ago

C final answer

upvoted 3 times

🗳️ 👤 **anandkl80** 3 years, 8 months ago

Answer B: Oracle Data Pump technology enables very high-speed movement of data and metadata from one database to another. Oracle Data Pump is available only on Oracle Database 10g release 1 (10.1) and later. ... Moving Data Between Different Database Versions. Original Export and Import Versus Data Pump Export and Import.

upvoted 1 times

🗳️ 👤 **Huy** 3 years, 7 months ago

Option B could be correct if it didn't mention 'Stop the application for the entire process'

upvoted 3 times

🗳️ 👤 **sellu** 3 years, 8 months ago

Option C is correct.

Option B is incorrect as it says application needs to be stopped till entire migration got completed

upvoted 4 times

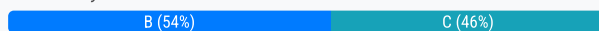
A company is using Amazon Aurora PostgreSQL for the backend of its application. The system users are complaining that the responses are slow. A database specialist has determined that the queries to Aurora take longer during peak times. With the Amazon RDS Performance Insights dashboard, the load in the chart for average active sessions is often above the line that denotes maximum CPU usage and the wait state shows that most wait events are IO:XactSync.

What should the company do to resolve these performance issues?

- A. Add an Aurora Replica to scale the read traffic.
- B. Scale up the DB instance class.
- C. Modify applications to commit transactions in batches.
- D. Modify applications to avoid conflicts by taking locks.

Suggested Answer: A

Community vote distribution



Huy Highly Voted 3 years, 7 months ago

C is answer. <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Reference.html>. Check frank's link.
upvoted 12 times

frankzeng Highly Voted 3 years, 7 months ago

C
<https://blog.dbi-services.com/aws-aurora-xactsync-batch-commit/>
upvoted 11 times

Hisayuki Most Recent 1 year, 4 months ago

Selected Answer: B
IO:XactSync with high CPU wait - means DB load exceed allocated vCPUs, so you should reduce those workloads or scale up to higher CPUs
But if the IO:XactSync is due to high commit, then you should modify your application to commit transactions in batches
upvoted 1 times

jitesh_k 1 year, 5 months ago

Answer seems C since IO:Xactsync is about commit/rollback.
<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/apg-waits.xactsync.html>
Batch commit seems to be the answer.
upvoted 1 times

jitesh_k 1 year, 5 months ago

Why not A? Add Aurora replica.
upvoted 3 times

Pranava_GCP 1 year, 9 months ago

Selected Answer: B
B. Scale up the DB instance class.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/apg-waits.xactsync.html#apg-waits.xactsync.actions.scalecpu>
upvoted 1 times

IhorK 1 year, 9 months ago

According to
<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/apg-waits.xactsync.html>
we should
- Scale up the CPU;
- Reduce the number of commits.
There are 2 correct answers: B and C. I don't know who to give preference to :(
upvoted 2 times

[Removed] 1 year, 11 months ago

Selected Answer: B

Hard choice between B&C, but i select B.

AWS recommends in documentation (<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/apg-waits.xactsync.html>) to Scale up the CPU or Reduce the number of commits.

Modify application will require some time -> therefore Scaling up will help to resolve current issues with CPU

upvoted 1 times

  **Zdujgfr567783ff** 1 year, 11 months ago

Selected Answer: C

By default in PostgreSQL, single-row inserts auto-commit for each row. This means that the database must wait for durability (write to storage) for every insert. To improve performance, PostgreSQL supports multi-row inserts and disabling auto-commit.

upvoted 1 times

  **Zdujgfr567783ff** 1 year, 11 months ago

C

By default in PostgreSQL, single-row inserts auto-commit for each row. This means that the database must wait for durability (write to storage) for every insert. To improve performance, PostgreSQL supports multi-row inserts and disabling auto-commit.

upvoted 1 times

  **Pankaj24hrs** 2 years ago

B

If CPU would not be an issue then definitely C, since a high CPU then it effects writing to the Storage layer as well.

CPU pressure

A heavy workload might be preventing the Aurora storage daemon from getting sufficient CPU time.

To address CPU starvation issues, consider changing to an instance type with more CPU capacity

upvoted 1 times

  **Zdujgfr567783ff** 1 year, 11 months ago

why CPU?

IO:XactSync. PDF RSS. The IO:XactSync event occurs when the database is waiting for the Aurora storage subsystem to acknowledge the commit of a regular transaction, or the commit or rollback of a prepared transaction.

upvoted 1 times

  **Zdujgfr567783ff** 1 year, 11 months ago

By default in PostgreSQL, single-row inserts auto-commit for each row. This means that the database must wait for durability (write to storage) for every insert. To improve performance, PostgreSQL supports multi-row inserts and disabling auto-commit.

upvoted 1 times

  **guau** 2 years, 3 months ago

Selected Answer: C

After reading the comments C is the right answer

upvoted 2 times

  **anantarb** 2 years, 3 months ago

Selected Answer: C

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/apg-waits.xactsync.html#apg-waits.xactsync.actions.commits>

upvoted 3 times

  **Mintwater** 2 years, 2 months ago

vote for you!

C

upvoted 2 times

  **tucobbad** 2 years, 5 months ago

Selected Answer: B

(with typos fixed)

My answer is B. Despite we know that CPU is related to it as well as COMMIT frequency is high, tell me honestly what would be the fastest/easiest approach here to resolve the issue? Scale up the instance (B) or Modify an application to commit in batches (C).

Second, this post by Franck Pachot has plenty info which basically leds to CPU issue: <https://www.dbi-services.com/blog/aws-aurora-xactsync-batch-commit/>

upvoted 3 times

🗨️ 👤 **Sathish_dbs** 2 years, 4 months ago

There is no ask about fastest/easiest approach here. so as a company the best approach to be taken here which is fixing the performance issue in the code. you can't keep increasing the hardware forever without fixing the code first!

upvoted 2 times

🗨️ 👤 **tucobbad** 2 years, 5 months ago

Selected Answer: B

My answer is B. Despite we know that CPU is related to it as well as COMMIT frequency is high, tell me honestly what would be the fastest/easiest approach here to resolve the issue? Scale up the instance (B) or Modify an application to commit in batches (C).

Second, this post by Frank Pachot has plenty info which basically leads to CPU issue: <https://www.dbi-services.com/blog/aws-aurora-xactsync-batch-commit/>

upvoted 1 times

🗨️ 👤 **bikash7758** 2 years, 6 months ago

C IS CORRECT

upvoted 2 times

🗨️ 👤 **rags1482** 2 years, 7 months ago

Answer : C

upvoted 2 times

A database specialist deployed an Amazon RDS DB instance in Dev-VPC1 used by their development team. Dev-VPC1 has a peering connection with Dev-VPC2 that belongs to a different development team in the same department. The networking team confirmed that the routing between VPCs is correct; however, the database engineers in Dev-VPC2 are getting a timeout connections error when trying to connect to the database in Dev-VPC1.

What is likely causing the timeouts?

- A. The database is deployed in a VPC that is in a different Region.
- B. The database is deployed in a VPC that is in a different Availability Zone.
- C. The database is deployed with misconfigured security groups.
- D. The database is deployed with the wrong client connect timeout configuration.

Suggested Answer: C

Community vote distribution

C (100%)

  **manan728** Highly Voted 3 years, 7 months ago



C it is.

upvoted 6 times

  **Aesthet** Highly Voted 3 years, 7 months ago

C final answer

upvoted 6 times

  **Aesthet** 3 years, 7 months ago

"A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IP addresses. Instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, with a VPC in another AWS account, or with a VPC in a different AWS Region."

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.Scenarios.html

upvoted 5 times

  **Pranava_GCP** Most Recent 1 year, 9 months ago

Selected Answer: C

C. The database is deployed with misconfigured security groups.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/security-group-rules-reference.html#sg-rules-db-server>



upvoted 1 times

  **novice_expert** 3 years ago

Selected Answer: C

security group needs inbound rule

upvoted 3 times

  **tugboat** 3 years, 3 months ago

Selected Answer: C

Security groups need to be fixed

upvoted 2 times

  **Suresh108** 3 years, 6 months ago

CCCCC

upvoted 2 times

A company has a production environment running on Amazon RDS for SQL Server with an in-house web application as the front end. During the last application maintenance window, new functionality was added to the web application to enhance the reporting capabilities for management. Since the update, the application is slow to respond to some reporting queries.

How should the company identify the source of the problem?

- A. Install and configure Amazon CloudWatch Application Insights for Microsoft .NET and Microsoft SQL Server. Use a CloudWatch dashboard to identify the root cause.
- B. Enable RDS Performance Insights and determine which query is creating the problem. Request changes to the query to address the problem.
- C. Use AWS X-Ray deployed with Amazon RDS to track query system traces.
- D. Create a support request and work with AWS Support to identify the source of the issue.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **2025flakyt** Highly Voted 👍 3 years, 5 months ago

B is correct

Amazon RDS Performance Insights is a database performance tuning and monitoring feature that helps you quickly assess the load on your database, and determine when and where to take action. Performance Insights allows non-experts to detect performance problems with an easy-to-understand dashboard that visualizes database load.

<https://aws.amazon.com/rds/performance-insights/>

upvoted 7 times

🗳️ 👤 **chen0305_099** Most Recent 🕒 1 year, 8 months ago

🗳️ 🗳️ 🗳️ 🗳️

upvoted 2 times

🗳️ 👤 **Pranava_GCP** 1 year, 9 months ago

Selected Answer: B

B. Enable RDS Performance Insights and determine which query is creating the problem. Request changes to the query to address the problem.

"Performance Insights uses lightweight data collection methods that don't impact the performance of your applications, and makes it easy to see which SQL statements are causing the load, and why. It requires no configuration or maintenance, and is currently available for Amazon Aurora (PostgreSQL- and MySQL-compatible editions), Amazon RDS for PostgreSQL, MySQL, MariaDB, SQL Server and Oracle."

"Whether your database performance problem is due to database configuration or application design issues, you can quickly identify the bottleneck and see which SQL statements are contributing to it."

<https://aws.amazon.com/rds/performance-insights/>

upvoted 1 times

🗳️ 👤 **IhorK** 1 year, 9 months ago

Selected Answer: B

Performance Insights to analyze database performance by waits, SQL statements, hosts, or users.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PerfInsights.ActiveSessions.html

Amazon CloudWatch Application Insights for Microsoft .NET and Microsoft SQL Server allows you to monitor the health of the SQL server, general issues not related to a specific query. Metrics that are monitored for the SQL Server instance:

- CPUUtilization
- SQLServer:Buffer Manager cache hit ratio
- SQLServer:Locks Number of Deadlocks/sec
- VolumeQueueLength

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/appinsights-what-is.html>

upvoted 2 times

🗳️ 👤 **teo2157** 2 years, 4 months ago

Selected Answer: B

I think it can't be A because it's an in-house web application, saying that going for B
upvoted 1 times

  **niteshdba** 2 years, 8 months ago

i think it should be A: as an Application, insight can provide more details to SQL server
upvoted 2 times

  **sachin** 2 years, 11 months ago

I think this should be A.

The question nowhere states the slowness is due to query. it states response is slow to certain request for reporting.



CloudWatch Application Insights is For common problems in .NET and SQL application stacks, such as application latency, SQL Server failed backups, memory leaks, large HTTP requests, and canceled I/O operations, it provides additional insights that point to a possible root cause and steps for resolution. Built-in integration with AWS SSM OpsCenter allows you to resolve issues by running the relevant Systems Manager Automation document.

First we need to understand the reason of slow response. It could be application issue rather SQL issue where Performance Insight can be helpful
upvoted 3 times

  **novice_expert** 3 years ago



Selected Answer: B

additional features to improve management reporting capabilities => new queries needing Performance Insights to check
upvoted 2 times

  **tugboat** 3 years, 3 months ago

Selected Answer: B

New reports = new queries needing Performance Insights to check
upvoted 3 times

  **Sp230** 3 years, 6 months ago

A

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/cloudwatch-application-insights.html>

" For common problems in .NET and SQL application stacks, such as application latency, SQL Server failed backups, memory leaks, large HTTP requests, and canceled I/O operations, it provides additional insights that point to a possible root cause and steps for resolution"
upvoted 3 times

  **jove** 3 years, 5 months ago

Correct answer is B.. The issue is not about the infrastructure, it is about the additional features added during with the application maintenance.
upvoted 2 times

  **Scunningham99** 3 years, 6 months ago

New functionality added to update reporting capabilities. Reports from application are generated via queries. These queries need to be investigated via answer which is > B
upvoted 2 times

  **sachin** 2 years, 11 months ago

Reports are generated by SQL but nowhere it is mentioned that reporting sql is slow. It says response of reporting request is slow which could be due to application issue as well
upvoted 1 times

  **ChauPhan** 3 years, 6 months ago

B is the answer
upvoted 2 times

  **AM** 3 years, 7 months ago

sql server application issue. Does not say queries. A is the answer.
upvoted 2 times



  **ExtHo** 3 years, 7 months ago

Problem started due new functionality was added to the web application to enhance the reporting capabilities for management. Since the update, the application is slow to respond to some reporting queries so issue is with query performance not with application that points B as correct answer
upvoted 1 times

  **Suresh108** 3 years, 7 months ago

BBBBBBBBB

upvoted 2 times

  **Aesthet** 3 years, 8 months ago

B final answer

upvoted 3 times

  **swarndeep** 3 years, 8 months ago

Option B

upvoted 2 times

  **shantest1** 3 years, 8 months ago

B. Answer

upvoted 4 times

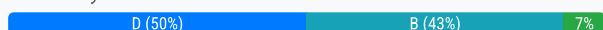
An electric utility company wants to store power plant sensor data in an Amazon DynamoDB table. The utility company has over 100 power plants and each power plant has over 200 sensors that send data every 2 seconds. The sensor data includes time with milliseconds precision, a value, and a fault attribute if the sensor is malfunctioning. Power plants are identified by a globally unique identifier. Sensors are identified by a unique identifier within each power plant. A database specialist needs to design the table to support an efficient method of finding all faulty sensors within a given power plant.

Which schema should the database specialist use when creating the DynamoDB table to achieve the fastest query time when looking for faulty sensors?

- A. Use the plant identifier as the partition key and the measurement time as the sort key. Create a global secondary index (GSI) with the plant identifier as the partition key and the fault attribute as the sort key.
- B. Create a composite of the plant identifier and sensor identifier as the partition key. Use the measurement time as the sort key. Create a local secondary index (LSI) on the fault attribute.
- C. Create a composite of the plant identifier and sensor identifier as the partition key. Use the measurement time as the sort key. Create a global secondary index (GSI) with the plant identifier as the partition key and the fault attribute as the sort key.
- D. Use the plant identifier as the partition key and the sensor identifier as the sort key. Create a local secondary index (LSI) on the fault attribute.

Suggested Answer: B

Community vote distribution



[Removed] 3 years, 8 months ago

I am going with D. As I understand it, you can have an item as a partition key and an item as a sort key to make a composite key. However, you cannot have two items as a partition key and a third item as sort key to make a composite key.

upvoted 18 times

awsjjj 2 years, 7 months ago

you can combine more than one attribute to create a partition key .

<https://aws.amazon.com/blogs/database/choosing-the-right-dynamodb-partition-key/>

"Use composite attributes. Try to combine more than one attribute to form a unique key, if that meets your access pattern. For example, consider an orders table with customerid#productid#countrycode as the partition key and order_date as the sort key, where the symbol # is used to split different field.

upvoted 3 times

FooBarBazBazinga 2 years ago

The access pattern for this question is: finding all faulty sensors within a given power plant.

Partition key DOESN'T support using condition expressions - you can only use them in sort key and in filter expressions. If you combine plantID and sensorID in a partition key, how do you find ALL faulty sensors within a single power plant? You simply cannot. With such a partition key, you always have to query for a specific sensor located in a specific power plant.

If 1 of the answers would be to place plantID and sensorID in sort key, then we could consider this option.

I will go with D.

upvoted 2 times

Jiang_aws1 2 years, 6 months ago

DynamoDB don't allow us do that way. we have to merge / combine into as 1 attribute

So B, C are wrong

upvoted 2 times

Jaypdv 3 years, 8 months ago

C. Putting the Fault attribute in a sparse index, based on <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-indexes-general-sparse-indexes.html>

upvoted 8 times

🗨️ 👤 **shantest1** 3 years, 8 months ago

Agreed C. Answer

upvoted 1 times

🗨️ 👤 **Hisayuki** Most Recent 1 year, 4 months ago

Selected Answer: D

If the answer is B, you need 200 partitions for each sensor.

upvoted 1 times

🗨️ 👤 **Sathish_dbs** 1 year, 7 months ago

Selected Answer: B

D is not having option to filter measurement time as an sensor could have been faulty last month but might have been fixed all the recent entries no faulty, though the question explicitly asked "based on the recent measurement" we need to assume it, and option D don't have this field in its solution, hence it will become a problem

upvoted 1 times

🗨️ 👤 **omle2** 1 year, 8 months ago

C is correct Answer I think.

Details are below.

"Power plants are identified by a globally unique identifier. Sensors are identified by a unique identifier within each power plant. "

It means that 'sensor_id' is not globally unique.

"A database specialist needs to design the table to support an efficient method of finding all faulty sensors within a given power plant."

The database specialist knows the 'plant_id' before executing the query.

As a prerequisite knowledge

- Secondary indexes like LSI and GSI in DynamoDB are not unique.

- The data in a LSI is organized by the same partition key as the base table, but with a different sort key.

upvoted 2 times

🗨️ 👤 **omle2** 1 year, 8 months ago

- A : x

Key (pk : plant_id - sk : measurement_time)

GSI (pk : plant_id - sk : fault)

If the plant_id is known, it is easy to use 'fault' to find the failed sensor.

1. Query for items with a specific 'plant_id' and 'fault' set to true.

2. The 'sensor_id' can be obtained from the acquired item.

If the measurement_time is exactly the same time in plant, the data could be overwritten.

- B : x

Key (pk : plant_id#sensor_id - sk : measurement_time)

LSI (pk : plant_id#sensor_id - sk : fault)

If you don't know the 'sensor_id', you can't query it.

upvoted 1 times

🗨️ 👤 **omle2** 1 year, 8 months ago

- C : o

Key (pk : plant_id#sensor_id - sk : measurement_time)

GSI (pk : plant_id - sk : fault)

As shown in A, fault sensor_id can be obtained using GSI.

The measurement_time is unique with the combination of plant_id and sensor_id.

- D : x

Key (pk : plant_id - sk : sensor_id)

LSI (pk : plant_id - sk : fault)

The combination of the partition key value and sort key value for each item must be unique.

Only one item per 'plant_id' and 'sensor_id' can be stored.

upvoted 1 times

  **Germaneli** 1 year, 8 months ago

Selected Answer: B

x A is out: a GSI must have the same partition key as the base table.

B. The PK consisting of partition key+sort key is unique. An LSI on the fault attribute is possible. That would satisfy the query.

x C. a GSI must have the same partition key as the base table.

x D. The PK consisting of partition key+sort key is not unique.

upvoted 1 times

  **Germaneli** 1 year, 8 months ago

Sorry, wrong reasoning: An LSI has the same partition key as the base table, but a different sort key. A GSI can have an altogether different partition key...

That means, B has a wrong LSI, C is not unique, D has a wrong LSI.

Remains A...

upvoted 1 times

  **Pranava_GCP** 1 year, 8 months ago

Selected Answer: D

D is correct answer.



Composite primary key: This is a combination of partition key and sort key. So, B & C are ruled out. Partition key composes of one attribute only.

"Partition key: A simple primary key, composed of one attribute known as the partition key. Attributes in DynamoDB are similar in many ways to fields or columns in other database systems."

A is wrong because no info about sensors identifier.

<https://aws.amazon.com/blogs/database/choosing-the-right-dynamodb-partition-key/>

upvoted 1 times

  **IhorK** 1 year, 9 months ago

Selected Answer: D

Composite primary key: This is a combination of partition key and sort key. So, B & C are wrong.

A - no info about sensors identifier.

upvoted 2 times

  **Zdujgfr567783ff** 1 year, 11 months ago

Selected Answer: D

looks like D

there is nothing about the timestamp in the task

upvoted 1 times

  **aviathor** 2 years ago

Selected Answer: B

A: The GSI key is not unique

B: The Primary key is unique because measurement time is the sort key, and the LSI allows us to narrow down to the required data

C: The primary key is not unique

D: The primary key is not unique.

The combination of plant id and sensor id is not unique since there will be many measurements for that key. Measurement time must therefore be part of the key to make it unique.

upvoted 4 times

🗨️ 👤 **aqiao** 2 years, 1 month ago

Selected Answer: D

You can't use more than one attributes as a partition key

upvoted 2 times

🗨️ 👤 **ken_test1234** 2 years, 1 month ago

Selected Answer: D

<https://aws.amazon.com/blogs/database/choosing-the-right-dynamodb-partition-key/>

upvoted 1 times

🗨️ 👤 **backbencher2022** 2 years, 2 months ago

Selected Answer: B

B is the only option where an index would have sensor id as well. We need faulty sensor as query output and if sensor id is included in an index then only we are going to get faulty sensor details in the output. If we notice carefully, none of the options except B have sensor id used in index and we are assuming that we will be using an index to get data faster. So, correct choice is the one which has the query attribute in the index.

upvoted 1 times

🗨️ 👤 **im_not_robot** 2 years, 3 months ago

Selected Answer: D

A is feasible. But we can fulfill the requirement with simpler approach (which is D)

B is completely wrong because we can not use that LSI at all for the requirement. To use the LSI, we need to specify the plant identifier as the primary key, but the primary key now contains sensor identifier so we can not use LSI.

C is feasible similar to A

D is the most simpler set up.

upvoted 2 times

🗨️ 👤 **guau** 2 years, 3 months ago

Selected Answer: C

Confusing question but I think it should be C

upvoted 1 times

🗨️ 👤 **anantarb** 2 years, 3 months ago

Selected Answer: D

A,C does not make sense: there is no need to create GSI, as question says find faulty sensors by plant. So when you have plant identifier as primary key, you can find faulty sensors in that plant by using LSI.

B does not make sense: Question does not mention anything finding item by measurement time. So there is no point having measurement time as sort key.

D is correct.

upvoted 2 times

🗨️ 👤 **BaskerS** 2 years, 3 months ago

D is correct Answer i think

Composite key is a Partition key and Sort Key, and then we cannot have another Sort Key

But

D works which is Partition Key Plant Identifier and Sort Key as Sensor Identified, then LSI as Fault Attribute

upvoted 1 times

A company is releasing a new mobile game featuring a team play mode. As a group of mobile device users play together, an item containing their statuses is updated in an Amazon DynamoDB table. Periodically, the other users' devices read the latest statuses of their teammates from the table using the BatchGetItem operation.

Prior to launch, some testers submitted bug reports claiming that the status data they were seeing in the game was not up-to-date. The developers are unable to replicate this issue and have asked a database specialist for a recommendation.

Which recommendation would resolve this issue?

- A. Ensure the DynamoDB table is configured to be always consistent.
- B. Ensure the BatchGetItem operation is called with the ConsistentRead parameter set to false.
- C. Enable a stream on the DynamoDB table and subscribe each device to the stream to ensure all devices receive up-to-date status information.
- D. Ensure the BatchGetItem operation is called with the ConsistentRead parameter set to true.

Suggested Answer: C

Community vote distribution



shantest1 Highly Voted 3 years, 7 months ago

D. answer
upvoted 15 times

ryuhei Highly Voted 2 years, 11 months ago

Selected Answer: D

Answer: D

https://docs.aws.amazon.com/ja_jp/amazondynamodb/latest/developerguide/HowItWorks.ReadConsistency.html

upvoted 6 times

Pranava_GCP Most Recent 1 year, 8 months ago

Selected Answer: D

D. Ensure the BatchGetItem operation is called with the ConsistentRead parameter set to true.

upvoted 2 times

IhorK 1 year, 9 months ago

Selected Answer: D

By default, BatchGetItem performs eventually consistent reads on every table in the request. If you want strongly consistent reads instead, you can set ConsistentRead to true for any or all tables.

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

upvoted 1 times

backbencher2022 2 years, 2 months ago

Selected Answer: D

D is correct

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

upvoted 2 times

Mintwater 2 years, 2 months ago

D.

By default, BatchGetItem performs eventually consistent reads on every table in the request. If you want strongly consistent reads instead, you can set ConsistentRead to true for any or all tables.

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

upvoted 2 times

sachin 2 years, 11 months ago

D is correct.

upvoted 2 times

elf78 2 years, 11 months ago



C is the answer. Questions states that "...periodically reads" which mean re using the BatchGetItem with consistent read (D) is not going to solve the problem since data might be already stale on the next read. However, by using Streams each device will be notified immediately.

upvoted 1 times

  **shammous** 2 years, 9 months ago

Consistent means getting the latest data. You are talking about the "eventual" case. Streams is usually used to trigger actions based on events happening on the db. D is probably the answer.

upvoted 1 times

  **dbaroger** 2 years, 12 months ago



Selected Answer: C

C is the most effective. Streams always will show the last update of the item, so, all users will read the same item status.

C would resolve the problem as well, but will double the read capacity unit and the cost will double too

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

upvoted 1 times

  **dbaroger** 2 years, 12 months ago

I meant "D" would double the RCUs.



upvoted 1 times

  **novice_expert** 3 years ago

Selected Answer: D

game's status data was out of current means query got eventual consistent data

upvoted 2 times


  **Dantas** 3 years, 2 months ago

Selected Answer: D

By default, BatchGetItem performs eventually consistent reads on every table in the request. If you want strongly consistent reads instead, you can set ConsistentRead to true for any or all tables.

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

upvoted 4 times

  **kped21** 3 years, 3 months ago

D, BatchGetItem with ConsistentRead

upvoted 1 times

  **user0001** 3 years, 3 months ago

D

By default, BatchGetItem performs eventually consistent reads on every table in the request. If you want strongly consistent reads instead, you can set ConsistentRead to true for any or all tables.

upvoted 1 times

  **mnzsql365** 3 years, 5 months ago

D



By default, BatchGetItem performs eventually consistent reads on every table in the request. If you want strongly consistent reads instead, you can set ConsistentRead to true for any or all tables.

upvoted 2 times

  **Suresh108** 3 years, 7 months ago

DDDDDDDD

upvoted 1 times

  **gdtypk** 3 years, 7 months ago

https://docs.aws.amazon.com/ja_jp/amazondynamodb/latest/developerguide/API_BatchGetItem_v20111205.html

upvoted 1 times

  **Aesthet** 3 years, 7 months ago

D final answer

upvoted 1 times

A company is running an Amazon RDS for MySQL Multi-AZ DB instance for a business-critical workload. RDS encryption for the DB instance is disabled. A recent security audit concluded that all business-critical applications must encrypt data at rest. The company has asked its database specialist to formulate a plan to accomplish this for the DB instance.

Which process should the database specialist recommend?

- A. Create an encrypted snapshot of the unencrypted DB instance. Copy the encrypted snapshot to Amazon S3. Restore the DB instance from the encrypted snapshot using Amazon S3.
- B. Create a new RDS for MySQL DB instance with encryption enabled. Restore the unencrypted snapshot to this DB instance.
- C. Create a snapshot of the unencrypted DB instance. Create an encrypted copy of the snapshot. Restore the DB instance from the encrypted snapshot.
- D. Temporarily shut down the unencrypted DB instance. Enable AWS KMS encryption in the AWS Management Console using an AWS managed CMK. Restart the DB instance in an encrypted state.

Suggested Answer: A

Reference:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html>

Community vote distribution

C (100%)

🗳️ 👤 **Jaypdr** Highly Voted 3 years, 7 months ago

C.

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html#Overview.Encryption.Limitations>

upvoted 14 times

🗳️ 👤 **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: C

C. Create a snapshot of the unencrypted DB instance. Create an encrypted copy of the snapshot. Restore the DB instance from the encrypted snapshot.

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html#Overview.Encryption.Limitations>

"you can create a snapshot of your DB instance, and then create an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot, and thus you have an encrypted copy of your original DB instance. "

upvoted 2 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: C

snapshot

-> encrypted copy of the snapshot

-> Restore the DB instance from the encrypted snapshot.

upvoted 3 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: C

Its definitely C

upvoted 3 times

🗳️ 👤 **Shunpin** 3 years, 5 months ago

Ans: C

In AWS console, you have to "migrate" unencrypted snapshot to encrypted one then performing database restoration.

upvoted 1 times

🗳️ 👤 **jove** 3 years, 6 months ago

Selected Answer: C

It is C

upvoted 3 times

🗨️ 👤 **GMartinelli** 3 years, 6 months ago

Selected Answer: C

Option C

upvoted 4 times

🗨️ 👤 **guru_ji** 3 years, 6 months ago

Answer: C

upvoted 1 times

🗨️ 👤 **Suresh108** 3 years, 7 months ago

CCCCCCCC

upvoted 1 times

🗨️ 👤 **shantest1** 3 years, 8 months ago

C. Answer

upvoted 2 times

🗨️ 👤 **shantest1** 3 years, 8 months ago

A and D is incorrect

B is for Aurora Database not for RDS Database.

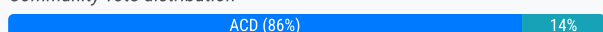
upvoted 2 times

A company is migrating its on-premises database workloads to the AWS Cloud. A database specialist performing the move has chosen AWS DMS to migrate an Oracle database with a large table to Amazon RDS. The database specialist notices that AWS DMS is taking significant time to migrate the data. Which actions would improve the data migration speed? (Choose three.)

- A. Create multiple AWS DMS tasks to migrate the large table.
- B. Configure the AWS DMS replication instance with Multi-AZ.
- C. Increase the capacity of the AWS DMS replication server.
- D. Establish an AWS Direct Connect connection between the on-premises data center and AWS.
- E. Enable an Amazon RDS Multi-AZ configuration.
- F. Enable full large binary object (LOB) mode to migrate all LOB data for all large tables.

Suggested Answer: CDE

Community vote distribution



Jaypdv Highly Voted 3 years, 8 months ago

ACD. Answer

Selecting A. based on https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html#CHAP_BestPractices.LargeTables

C. and D. are the only options that have to do with increasing performance, the others are irrelevant

upvoted 11 times

johnconnor 3 years, 6 months ago

it is going to take you days/weeks to establish a direct connection, if it is not from the start, it should not be an option

upvoted 3 times

Justu 3 years, 5 months ago

It should be not an option, but these three (ACD) are the means which fasten the migration. Other options wouldn't do that.

upvoted 2 times

guru_ji Highly Voted 3 years, 7 months ago

Correct Answer: ACF

Full LOB mode – migrates all LOB data, piecewise in chunks (you provide LOB chunk size)

upvoted 6 times

tcl08 Most Recent 1 year, 3 months ago

D can't be part of the answer. It takes about a week or more to setup a direct connect connection. I will think it ACF

upvoted 1 times

Pranava_GCP 1 year, 8 months ago

Selected Answer: ACD

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html

B & E. Multi AZ are for improving availability should a storage issue occur, not contributing to data migration speed so B&E are out. F is opposite, "Full LOB mode provides the convenience of moving all LOB data in your tables, but the process can have a significant impact on performance." so F is out.

upvoted 1 times

cloudbusting 1 year, 12 months ago

it should be ACF

<https://repost.aws/knowledge-center/dms-improve-speed-lob-data>

upvoted 1 times

MrAliMohsan 2 years ago

Selected Answer: ACD

The answer I was searching in the comments, I will post it here.

For those searching for a reason to select D, and how it improves migration speed.

AWS Direct Connect can reduce network costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections

upvoted 1 times

🗨️ **aviathor** 2 years ago

Selected Answer: ACD

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html#CHAP_BestPractices.Performance

upvoted 1 times

🗨️ **aqiao** 2 years, 1 month ago

F is wrong: In full LOB mode AWS DMS migrates all LOBs from source to target regardless of size. In this configuration, AWS DMS has no information about the maximum size of LOBs to expect. Thus, LOBs are migrated one at a time, piece by piece. Full LOB mode can be quite slow.

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Tasks.LOBSupport.html

upvoted 2 times

🗨️ **sk1974** 2 years, 3 months ago

How can we break a single table into multiple DMS tasks .If the DB had multiple tables , they can be split into multiple DMS tasks . So , we cannot assume that 'large table' to mean to multiple tables. So , I will go with CDE

upvoted 1 times

🗨️ **aviathor** 2 years ago

"To do this, split the table into segments and load the segments in parallel in the same migration task."

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html#CHAP_BestPractices.Performance

upvoted 1 times

🗨️ **rags1482** 2 years, 7 months ago

With Amazon RDS databases, it's a good idea to turn off backups and Multi-AZ until the cutover.

So B and E eliminated

F - Full LOB mode impacts performance

So ACD right answer

upvoted 1 times

🗨️ **Satprave** 2 years, 7 months ago

F - Full LOB mode downgrade the performance

Full LOB mode migrates all LOB data in your tables, regardless of size. Full LOB mode provides the convenience of moving all LOB data in your tables, but the process can have a significant impact on performance.

upvoted 2 times

🗨️ **Dantas** 2 years, 12 months ago

Selected Answer: ACD

"To improve the performance when migrating a large table, break the migration into more than one task. To break the migration into multiple tasks using row filtering, use a key or a partition key" (A)

A number of factors affect the performance of your AWS DMS migration:

- Resource availability on the source.
- The available network throughput. (D)
- The resource capacity of the replication server. (C)
- The ability of the target to ingest changes.
- The type and distribution of source data.
- The number of objects to be migrated.

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html#CHAP_BestPractices.LargeTables

upvoted 5 times

🗨️ **whoareyou** 3 years ago

Selected Answer: ACF

It takes much to deploy the direct connect, so D is not the correct answer.

I choose ACF

upvoted 2 times

  **KaranGandhi30** 3 years ago

Selected Answer: ACD

B & E. Multi AZ doesn't increase speed it will add more time.

F. LOB mode will ensure complete movement but not faster


upvoted 2 times

  **novice_expert** 3 years ago

Selected Answer: ACD

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_BestPractices.html#CHAP_BestPractices.LargeTables

upvoted 2 times

  **awsmonster** 3 years, 3 months ago

ACF


F:

<https://aws.amazon.com/premiumsupport/knowledge-center/dms-improve-speed-lob-data/>

If you have LOBs that are larger than a few megabytes, then you can create a separate AWS DMS task with Full LOB mode. It's a best practice to create the separate task on a new replication instance to migrate these tables alone. Option A stated to create multiple DMS tasks.

Having an expensive Direct Connect will not resolve this issue

upvoted 2 times

  **jove** 3 years, 5 months ago

Strange question. If you've already started the migration of a large table and observe that the DMS is slow it's too late for making such suggested changes. If you're planning to migrate a large table then go with A,C and D.

upvoted 1 times

A company is migrating a mission-critical 2-TB Oracle database from on premises to Amazon Aurora. The cost for the database migration must be kept to a minimum, and both the on-premises Oracle database and the Aurora DB cluster must remain open for write traffic until the company is ready to completely cut over to Aurora.

Which combination of actions should a database specialist take to accomplish this migration as quickly as possible? (Choose two.)

- A. Use the AWS Schema Conversion Tool (AWS SCT) to convert the source database schema. Then restore the converted schema to the target Aurora DB cluster.
- B. Use Oracle's Data Pump tool to export a copy of the source database schema and manually edit the schema in a text editor to make it compatible with Aurora.
- C. Create an AWS DMS task to migrate data from the Oracle database to the Aurora DB cluster. Select the migration type to replicate ongoing changes to keep the source and target databases in sync until the company is ready to move all user traffic to the Aurora DB cluster.
- D. Create an AWS DMS task to migrate data from the Oracle database to the Aurora DB cluster. Once the initial load is complete, create an AWS Kinesis Data Firehose stream to perform change data capture (CDC) until the company is ready to move all user traffic to the Aurora DB cluster.
- E. Create an AWS Glue job and related resources to migrate data from the Oracle database to the Aurora DB cluster. Once the initial load is complete, create an AWS DMS task to perform change data capture (CDC) until the company is ready to move all user traffic to the Aurora DB cluster.

Suggested Answer: AD



Community vote distribution

AC (100%)

  **Jaypdy** Highly Voted 3 years, 7 months ago

AC. Answer

upvoted 13 times

  **Suresh108** Highly Voted 3 years, 7 months ago

AC --- locked

A - OK



B - NOT OK Text editor

C - ok though it doesn't talk about full load

D- NOT ok firehose

E - not ok Glue

upvoted 6 times

  **guru_ji** 3 years, 7 months ago

I agree with you.

upvoted 1 times

  **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: AC

AC are correct answer

SCT + DMS

upvoted 2 times

  **FooBarBazBazinga** 2 years ago

Selected Answer: AC

AC. asnwer

upvoted 1 times

  **backbencher2022** 2 years, 2 months ago

Selected Answer: AC

A&C for sure

upvoted 1 times

  **ninjalight25** 2 years, 2 months ago

Selected Answer: AC

AC --- locked

upvoted 1 times

  **novice_expert** 3 years ago

Selected Answer: AC

A. Use the AWS Schema Conversion Tool (AWS SCT) to convert the source database schema. Then restore the converted schema to the target Aurora DB cluster.

x B. edit

C. Create an AWS DMS task to migrate data from the Oracle database to the Aurora DB cluster. Select the migration type to replicate ongoing changes to keep the source and target databases in sync until the company is ready to move all user traffic to the Aurora DB cluster.

x D. kinesis firehose



x E. Glue job

upvoted 2 times

  **Scunningham99** 3 years, 6 months ago

a and c all day long



upvoted 2 times

  **guru_ji** 3 years, 7 months ago

I got this Question in exam.

60% questions came in actual exam from this 145 set. Bunch of new Questions.

upvoted 1 times

  **Aesthet** 3 years, 7 months ago

AC final answer

upvoted 3 times

  **agrawalachin** 3 years, 7 months ago

A C is the answer

upvoted 2 times

A company has a 20 TB production Amazon Aurora DB cluster. The company runs a large batch job overnight to load data into the Aurora DB cluster. To ensure the company's development team has the most up-to-date data for testing, a copy of the DB cluster must be available in the shortest possible time after the batch job completes.

How should this be accomplished?

- A. Use the AWS CLI to schedule a manual snapshot of the DB cluster. Restore the snapshot to a new DB cluster using the AWS CLI.
- B. Create a dump file from the DB cluster. Load the dump file into a new DB cluster.
- C. Schedule a job to create a clone of the DB cluster at the end of the overnight batch process.
- D. Set up a new daily AWS DMS task that will use cloning and change data capture (CDC) on the DB cluster to copy the data to a new DB cluster. Set up a time for the AWS DMS stream to stop when the new cluster is current.

Suggested Answer: B

Community vote distribution

C (100%)

  **shantest1** Highly Voted 3 years, 7 months ago

C. answer

Aurora can do cloning.

upvoted 14 times

  **novice_expert** Highly Voted 3 years ago

Selected Answer: C

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Clone.html>

Creating a clone is faster and more space-efficient than physically copying the data using other techniques, such as restoring a snapshot.

upvoted 5 times

  **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: C

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Clone.html>

D. does not make sense, DMS doesn't use cloning.

upvoted 2 times

  **megadba** 3 years, 1 month ago

Selected Answer: C

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Clone.html>

upvoted 1 times

  **Balaji_Cloud** 3 years, 4 months ago

C - because Clone should kick-off after batch completes

upvoted 2 times

  **GMartinelli** 3 years, 5 months ago

Selected Answer: C



Option C, Clone is easy and fast

upvoted 2 times

  **Suresh108** 3 years, 6 months ago

CCCC - faster clones

upvoted 1 times

  **Aesthet** 3 years, 7 months ago

C final answer

upvoted 2 times

  **manan728** 3 years, 7 months ago

C is correct.
upvoted 1 times

A company has two separate AWS accounts: one for the business unit and another for corporate analytics. The company wants to replicate the business unit data stored in Amazon RDS for MySQL in us-east-1 to its corporate analytics Amazon Redshift environment in us-west-1. The company wants to use AWS DMS with Amazon RDS as the source endpoint and Amazon Redshift as the target endpoint. Which action will allow AWS DMS to perform the replication?

- A. Configure the AWS DMS replication instance in the same account and Region as Amazon Redshift.
- B. Configure the AWS DMS replication instance in the same account as Amazon Redshift and in the same Region as Amazon RDS.
- C. Configure the AWS DMS replication instance in its own account and in the same Region as Amazon Redshift.
- D. Configure the AWS DMS replication instance in the same account and Region as Amazon RDS.

Suggested Answer: C

Reference:



<https://aws.amazon.com/premiumsupport/knowledge-center/dms-redshift-target-endpoint/>

Community vote distribution

A (100%)

 **Jaypdr**  3 years, 7 months ago

Sorry I meant A. https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.Redshift.html
upvoted 17 times

 **58a2d17**  1 year, 1 month ago

A. not C. Ref says it all : your Amazon Redshift cluster must be in the same account and same AWS Region as the replication instance
upvoted 1 times

 **Pranava_GCP** 1 year, 8 months ago

Selected Answer: A

A. Configure the AWS DMS replication instance in the same account and Region as Amazon Redshift.

"The Amazon Redshift cluster must be in the same AWS account and same AWS Region as the replication instance."

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.Redshift.html

upvoted 1 times

 **Zdujgr567783ff** 1 year, 11 months ago

Selected Answer: A

use target side
upvoted 2 times

 **khun** 2 years, 5 months ago

Selected Answer: A

Ans A.

The Amazon Redshift cluster must be in the same AWS account and the same AWS Region as the replication instance.

upvoted 2 times

 **novice_expert** 3 years ago

Selected Answer: A

A. Configure the AWS DMS replication instance in the same account and Region as Target Database
upvoted 4 times

 **peacegrace** 3 years, 3 months ago

Selected Answer: A

A is the answer. Refer : https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.Redshift.html
upvoted 2 times

 **Suresh108** 3 years, 6 months ago

AAA - locked -

upvoted 1 times

🗨️ 👤 **Aesthet** 3 years, 7 months ago

A final answer

"The Amazon Redshift cluster must be in the same AWS account and same AWS Region as the replication instance."

upvoted 4 times

🗨️ 👤 **manan728** 3 years, 7 months ago

A is right.

upvoted 1 times

🗨️ 👤 **shantest1** 3 years, 7 months ago

A. Answer

upvoted 2 times

🗨️ 👤 **Jaypdy** 3 years, 7 months ago

B. Answer

Ref https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.Redshift.html

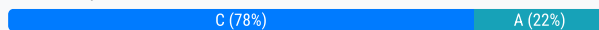
upvoted 2 times

A database specialist is managing an application in the us-west-1 Region and wants to set up disaster recovery in the us-east-1 Region. The Amazon Aurora MySQL DB cluster needs an RPO of 1 minute and an RTO of 2 minutes. Which approach meets these requirements with no negative performance impact?

- A. Enable synchronous replication.
- B. Enable asynchronous binlog replication.
- C. Create an Aurora Global Database.
- D. Copy Aurora incremental snapshots to the us-east-1 Region.

Suggested Answer: A

Community vote distribution



Jaypdv Highly Voted 3 years, 7 months ago

C.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-global-database-disaster-recovery.html>
upvoted 13 times

Pranava_GCP Most Recent 1 year, 8 months ago

Selected Answer: C

C. Create an Aurora Global Database.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-global-database-disaster-recovery.html>
upvoted 2 times

ninjalight25 2 years, 2 months ago

Selected Answer: A

Enable synchronous replication: Synchronous replication ensures that data is written to both the primary and standby DB clusters before the transaction is committed. This means that the secondary cluster is always in sync with the primary, with minimal lag time, and can be promoted to primary in case of a failure. This approach meets the RPO and RTO requirements and ensures there is no data loss or downtime.

upvoted 2 times

jitesh_k 1 year, 5 months ago

How about performance impact?

upvoted 1 times

novice_expert 3 years ago

Selected Answer: C

C. Create an Aurora Global Database.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-global-database-disaster-recovery.html>

Managed planned failover RPO = 0 sec, RTO = minutes

upvoted 4 times

AriraAWS 3 years, 3 months ago

Selected Answer: C

C is the right answer.

upvoted 1 times

Aesthet 3 years, 6 months ago

C final answer

upvoted 3 times

shantest1 3 years, 8 months ago



C. I believe

upvoted 2 times

  **shantest1** 3 years, 7 months ago



Sync replication not available cross region If I am not wrong.

upvoted 2 times

  **Hariru** 3 years, 3 months ago

So if we want to have it in the same region, synchronous replication would be a better solution?

upvoted 1 times

  **khchan123** 3 years ago

No. You can use multi-AZ read replica for single region DR.

upvoted 1 times

A gaming company is developing a new mobile game and decides to store the data for each user in Amazon DynamoDB. To make the registration process as easy as possible, users can log in with their existing Facebook or Amazon accounts. The company expects more than 10,000 users. How should a database specialist implement access control with the LEAST operational effort?

- A. Use web identity federation on the mobile app and AWS STS with an attached IAM role to get temporary credentials to access DynamoDB.
- B. Use web identity federation on the mobile app and create individual IAM users with credentials to access DynamoDB.
- C. Use a self-developed user management system on the mobile app that lets users access the data from DynamoDB through an API.
- D. Use a single IAM user on the mobile app to access DynamoDB.

Suggested Answer: A

Reference:

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


Community vote distribution

A (100%)

 **shantest1** Highly Voted 3 years, 8 months ago

A. answer

upvoted 11 times

 **chikorita** Most Recent 1 year, 7 months ago

right in front of my salad Amazon Cognito exists

upvoted 1 times

 **Pranava_GCP** 1 year, 8 months ago

Selected Answer: A

A. Use web identity federation on the mobile app and AWS STS with an attached IAM role to get temporary credentials to access DynamoDB.

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

"If you are writing an application targeted at large numbers of users, you can optionally use web identity federation for authentication and authorization. Web identity federation removes the need for creating individual users. Instead, users can sign in to an identity provider and then obtain temporary security credentials from AWS Security Token Service (AWS STS). The app can then use these credentials to access AWS services.

Web identity federation supports the following identity providers:

Login with Amazon

Facebook

Google"

upvoted 1 times

 **khun** 2 years, 5 months ago

Selected Answer: A

A, Web identity Fed > STS > IAM > tempo creds

upvoted 1 times

 **elf78** 2 years, 11 months ago

Selected Answer: A

IAM Role is always preferred method.


upvoted 2 times

 **novice_expert** 3 years ago

Selected Answer: A

Use web identity federation on the mobile app and AWS STS with an attached IAM role to get temporary credentials to access DynamoDB.

upvoted 1 times

 **tugboat** 3 years, 3 months ago

Selected Answer: A



obvious

upvoted 1 times

  **Suresh108** 3 years, 6 months ago



AAAAAAAAA

upvoted 2 times

  **Aesthet** 3 years, 6 months ago

A final answer

upvoted 2 times

  **agrawalachin** 3 years, 7 months ago

Answer is A - least operational

upvoted 1 times

A large retail company recently migrated its three-tier ecommerce applications to AWS. The company's backend database is hosted on Amazon Aurora PostgreSQL. During peak times, users complain about longer page load times. A database specialist reviewed Amazon RDS Performance Insights and found a spike in IO:XactSync wait events. The SQL attached to the wait events are all single INSERT statements. How should this issue be resolved?

- A. Modify the application to commit transactions in batches
- B. Add a new Aurora Replica to the Aurora DB cluster.
- C. Add an Amazon ElastiCache for Redis cluster and change the application to write through.
- D. Change the Aurora DB cluster storage to Provisioned IOPS (PIOPS).

Suggested Answer: B

Community vote distribution

A (100%)

 **shantest1** Highly Voted 3 years, 8 months ago

A. answer
upvoted 14 times

 **gelsm** 3 years, 6 months ago

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Reference.html>

"This wait most often arises when there is a very high rate of commit activity on the system. You can sometimes alleviate this wait by modifying applications to commit transactions in batches. "
upvoted 4 times

 **RBSK** Highly Voted 2 years, 7 months ago

Selected Answer: A

D - There is no option to setup PIOPS for Aurora. Only RDS has it
upvoted 6 times

 **Mintwater** 2 years, 1 month ago

Thanks for your explanation for D
upvoted 2 times

 **adelcold** Most Recent 1 year, 11 months ago

Selected Answer: A


<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/apg-waits.xactsync.html>

Monitor your resources
Scale up the CPU
Increase network bandwidth
Reduce the number of commits
upvoted 1 times

 **novice_expert** 3 years ago

Selected Answer: A

avoid high rate of commits by using batch commit
upvoted 2 times

 **soyyodario** 3 years, 3 months ago

Selected Answer: A

Answer is A.

Actions

We recommend different actions depending on the causes of your wait event.

- 1.- Monitor your resources
- 2.- Scale up the CPU
- 3.- Increase network bandwidth
- 4.- Reduce the number of commits

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/apg-waits.xactsync.html>
upvoted 6 times

🗲️ 👤 **jove** 3 years, 5 months ago

Selected Answer: A

To reduce the number of commits, combine statements into transaction blocks.
upvoted 3 times

🗲️ 👤 **Justu** 3 years, 5 months ago

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/apg-waits.xactsync.html>
upvoted 1 times

🗲️ 👤 **Suresh108** 3 years, 7 months ago

AAAAAA

upvoted 1 times

🗲️ 👤 **Aesthet** 3 years, 7 months ago

A final answer

upvoted 1 times

🗲️ 👤 **agrawalachin** 3 years, 7 months ago

Yes, A is correct

upvoted 1 times

🗲️ 👤 **manan728** 3 years, 7 months ago

Yup A is right for this use case

upvoted 1 times

A company uses Amazon DynamoDB as the data store for its ecommerce website. The website receives little to no traffic at night, and the majority of the traffic occurs during the day. The traffic growth during peak hours is gradual and predictable on a daily basis, but it can be orders of magnitude higher than during off-peak hours.

The company initially provisioned capacity based on its average volume during the day without accounting for the variability in traffic patterns. However, the website is experiencing a significant amount of throttling during peak hours. The company wants to reduce the amount of throttling while minimizing costs.

What should a database specialist do to meet these requirements?

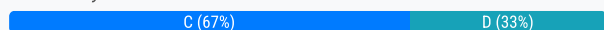
- A. Use reserved capacity. Set it to the capacity levels required for peak daytime throughput.
- B. Use provisioned capacity. Set it to the capacity levels required for peak daytime throughput.
- C. Use provisioned capacity. Create an AWS Application Auto Scaling policy to update capacity based on consumption.
- D. Use on-demand capacity.

Suggested Answer: D

Reference:

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

Community vote distribution



Sathish_dbs 1 year, 7 months ago

reduce cost - on-demand costlier

predictable, gradual increase - auto scaling

upvoted 2 times

leotoras 1 year, 10 months ago

C is correct based on documentation: When you create a DynamoDB table, auto scaling is the default capacity setting, but you can also enable auto scaling on any table that does not have it active. Behind the scenes, as illustrated in the following diagram, DynamoDB auto scaling uses a scaling policy in Application Auto Scaling.

<https://aws.amazon.com/blogs/database/amazon-dynamodb-auto-scaling-performance-and-cost-optimization-at-any-scale/>

upvoted 1 times

aviathor 1 year, 11 months ago

Selected Answer: D

If C had been worded differently, like "Provisioned capacity with auto-scaling", I might have chosen C. But "AWS Application Auto-Scaling" does not make sense in the DynamoDB context. Therefore I choose "On-Demand" which is very flexible in terms of throughput.

upvoted 1 times

guau 2 years, 3 months ago

CCCCCCCCCCCC

upvoted 1 times

Arun32 2 years, 6 months ago

C for me..

DynamoDB auto scaling uses a scaling policy in Application Auto Scaling. To configure auto scaling in DynamoDB, you set the minimum and maximum levels of read and write capacity in addition to the target utilization percentage. Auto scaling uses Amazon CloudWatch to monitor a table's read and write capacity metrics. To do so, it creates CloudWatch alarms that track consumed capacity.

upvoted 1 times

awsjji 2 years, 7 months ago

Selected Answer: D

no traffic at night and orders of magnitude high traffic in off peak. I will go with D

upvoted 2 times

kush_sumit 2 years, 10 months ago

Selected Answer: C

C: As the pattern is predictable load with magnitude of spikes this could be handled by autoscaling. As per AWS autoscaling is cost saving as compared to on-demand. If the pattern was unpredictable On-demand would be good

upvoted 1 times

🗨️ **sachin** 2 years, 11 months ago

but it also says "traffic increase during peak hours is steady and predictable" and the firm also wants to reduce the throttling " so it is either C or D I think ..

C is more likely as traffic pattern is predictable

upvoted 1 times

🗨️ **sachin** 2 years, 11 months ago

It states that there is a lot of throttling throughout the day and the firm wants to minimize the expenditure and cost. Provisioned is costlier and during night hours the traffic is almost nil.

So whatever small (average of day) capacity you are provisioning will be wasted at night.

I believe on-demand in this case is a better choice.

upvoted 1 times

🗨️ **Dantas** 2 years, 12 months ago

Selected Answer: C

"... traffic increase during peak hours is steady and predictable ..."

upvoted 1 times

🗨️ **novice_expert** 3 years ago

Selected Answer: C

- traffic increase during peak hours is steady and predictable => provisioned

- no uses in night. => on demand (but costly)

so provision some + auto scale with a target 70% utilization set

D will also work but costly

upvoted 3 times

🗨️ **mike3g2000** 3 years, 2 months ago

I go with C.

"To understand how DynamoDB auto scaling works, suppose that you have a table named ProductCatalog. The table is bulk-loaded with data infrequently, so it doesn't incur very much write activity. However, it does experience a high degree of read activity, which varies over time. By monitoring the Amazon CloudWatch metrics for ProductCatalog, you determine that the table requires 1,200 read capacity units (to avoid DynamoDB throttling read requests when activity is at its peak). You also determine that ProductCatalog requires 150 read capacity units at a minimum, when read traffic is at its lowest point.

Within the range of 150 to 1,200 read capacity units, you decide that a target utilization of 70 percent would be appropriate for the ProductCatalog table. Target utilization is the ratio of consumed capacity units to provisioned capacity units, expressed as a percentage. Application Auto Scaling uses its target tracking algorithm to ensure that the provisioned read capacity of ProductCatalog is adjusted as required so that utilization remains at or near 70 percent."

upvoted 2 times

🗨️ **RotterDam** 3 years, 2 months ago

Very tricky. Feels like (D)

Even if you enable Autoscaling - but you will still pay the same rate during off-peak hours. The provisioned capacity will never fall below the preconfigured one during night when there's no traffic. Plus the question WANTS to reduce expenditure - which means with OnDemand there's barely any charge during off-peak hours compared to fully charged if you use provisioned capacity using the baseline configured RCU/WCUs.

upvoted 3 times

🗨️ **whn** 3 years, 1 month ago

Remember for Provisioned with Auto-Scaling you are basically paying for throughput 24/7. Whereas for On-Demand Scaling you pay per request.

This means for applications still in development or low traffic applications, it might be more economical to use On-Demand Scaling and not worry about provisioning throughput. However, at scale, this can quickly shift once you have a more consistent usage pattern.

[https://dynobase.dev/dynamodb-on-demand-vs-provisioned-](https://dynobase.dev/dynamodb-on-demand-vs-provisioned-scaling/#:~:text=Remember%20for%20Provisioned%20with%20Auto,not%20worry%20about%20provisioning%20throughput.)

[scaling/#:~:text=Remember%20for%20Provisioned%20with%20Auto,not%20worry%20about%20provisioning%20throughput.](https://dynobase.dev/dynamodb-on-demand-vs-provisioned-scaling/#:~:text=Remember%20for%20Provisioned%20with%20Auto,not%20worry%20about%20provisioning%20throughput.)

upvoted 2 times

🗨️ **Paulv82003** 1 year, 11 months ago

The key is "The traffic growth during peak hours is gradual and predictable on a daily basis" provisioned with auto scaling

upvoted 1 times

🗉 👤 **user0001** 3 years, 2 months ago

C

Provisioned Mode

If you choose provisioned mode, you specify the number of reads and writes per second that you require for your application. You can use auto scaling to adjust your table's provisioned capacity automatically in response to traffic changes. This helps you govern your DynamoDB use to stay at or below a defined request rate in order to obtain cost predictability.

Provisioned mode is a good option if any of the following are true:

You have predictable application traffic.

You run applications whose traffic is consistent or ramps gradually.

You can forecast capacity requirements to control costs.

upvoted 4 times

🗉 👤 **tugboat** 3 years, 3 months ago

Selected Answer: C

They know peak hours

AWS Application Auto Scaling policy is good for it.

upvoted 1 times

🗉 👤 **Shinytopology** 3 years, 4 months ago

C. DynamoDB autoscaling saves costs comparing to on-demand (which costs a bit higher for the NoOps benefit.)

<https://aws.amazon.com/blogs/database/amazon-dynamodb-auto-scaling-performance-and-cost-optimization-at-any-scale/>

upvoted 2 times

🗉 👤 **wcx** 3 years, 4 months ago

"traffic increase during peak hours is steady and predictable"

C is the answer.

upvoted 1 times

A company uses an Amazon RDS for PostgreSQL DB instance for its customer relationship management (CRM) system. New compliance requirements specify that the database must be encrypted at rest.

Which action will meet these requirements?

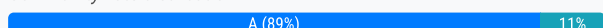
- A. Create an encrypted copy of manual snapshot of the DB instance. Restore a new DB instance from the encrypted snapshot.
- B. Modify the DB instance and enable encryption.
- C. Restore a DB instance from the most recent automated snapshot and enable encryption.
- D. Create an encrypted read replica of the DB instance. Promote the read replica to a standalone instance.

Suggested Answer: C

Reference:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html>

Community vote distribution



shantest1 Highly Voted 3 years, 8 months ago

A. Answer

upvoted 13 times

scottkerker 3 years, 5 months ago

This page has detailed steps for MySQL and Maria in terms of the encryption of an unencrypted RDS instance.

<https://aws.amazon.com/premiumsupport/knowledge-center/rds-encrypt-instance-mysql-mariadb/>

upvoted 1 times

toppic26 3 years, 6 months ago

From the reference: You can only enable encryption for an Amazon RDS DB instance when you create it, not after the DB instance is created.

However, because you can encrypt a copy of an unencrypted snapshot, you can effectively add encryption to an unencrypted DB instance. That is, you can create a snapshot of your DB instance, and then create an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot, and thus you have an encrypted copy of your original DB instance. For more information, see Copying a snapshot.

upvoted 5 times

Dantas 3 years, 4 months ago

Don't the words "manual snapshot" invalidate the answer "A"?

upvoted 1 times

Dantas 3 years, 2 months ago

A. Answer -> Answering to my own question: No! I've just tried to restore unencrypted manual and automatic snapshots into an encrypted db instance and it isn't allowed. If you want to launch an encrypted rds instance, you need to create an encrypted copy of the unencrypted snapshot.

upvoted 1 times

redman50 Most Recent 2 years, 2 months ago

Sign in to the AWS Management Console and navigate to the Amazon RDS dashboard.

Select the DB instance that you want to encrypt.

Click the "Modify" button.

In the "Encryption" section, select the option to "Enable encryption".

Choose the KMS encryption key that you want to use or create a new one.

Click "Continue" and review the summary of changes.

Click "Modify DB instance" to apply the changes.

Note that the encryption process will initiate a snapshot of the DB instance, encrypt it, and restore the encrypted data from the snapshot, so there will be a brief period of downtime while the encryption process is completed.

upvoted 1 times

megadba 2 years, 11 months ago

Selected Answer: D

D. is correct.

Create Read Replica encrypted enable and promote standalone instance.

A. The snapshot doesn't encrypted option.

B. Unencrypted instance is not enable encrypted.

C. Also automated snapshot is not enable encrypted.

upvoted 1 times

🗳️ 👤 **minhntm** 2 years, 11 months ago

Selected Answer: A

A. Answer

upvoted 1 times

🗳️ 👤 **niau** 2 years, 11 months ago

Selected Answer: A

You can only encrypt an Amazon RDS DB instance when you create it, not after the DB instance is created.

However, because you can encrypt a copy of an unencrypted snapshot, you can effectively add encryption to an unencrypted DB instance. That is, you can create a snapshot of your DB instance, and then create an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot, and thus you have an encrypted copy of your original DB instance

upvoted 1 times

🗳️ 👤 **praffuln** 3 years ago

A is correct, from the page DMS used for ongoing replication.

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/encrypt-an-existing-amazon-rds-for-postgresql-db-instance.html>

AWS DMS – You can use AWS Database Migration Service (AWS DMS) to replicate changes from the source DB to the target DB. It is important to keep the source and target DB in sync to keep downtime to a minimum. For information about setting up AWS DMS and creating tasks, see the AWS DMS documentation.

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: A

Manual Snapshot -> Create an encrypted copy -> Restore a new DB instance from the encrypted snapshot.

upvoted 1 times

🗳️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: A

correct

upvoted 2 times

🗳️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: A

Correct option

upvoted 1 times

🗳️ 👤 **kped21** 3 years, 3 months ago

A - is wrong, something has changed recently.

I took a snapshot and tried to copy and encrypt it, it does not allow unencrypted to encrypted.

The best option is C, C works as is take any snapshot or manual snapshot and restore to new encrypted cluster.

upvoted 1 times

🗳️ 👤 **Sisun** 3 years, 5 months ago

Selected Answer: A

A - correct

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/encrypt-an-existing-amazon-rds-for-postgresql-db-instance.html>

You can enable encryption for an Amazon RDS DB instance when you create it, but not after it's created. However, you can add encryption to an unencrypted DB instance by creating a snapshot of your DB instance, and then creating an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot to get an encrypted copy of your original DB instance. The pattern uses AWS Database Migration Service (AWS DMS) to migrate data and AWS Key Management Service (AWS KMS) for encryption.

upvoted 2 times

🗳️ 👤 **akiraklaus** 3 years, 5 months ago

All aswer is erro, is necessary utilization DMS



<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/encrypt-an-existing-amazon-rds-for-postgresql-db-instance.html>

upvoted 1 times

  **ChauPhan** 3 years, 6 months ago



A is correct

upvoted 1 times

  **Aesthet** 3 years, 7 months ago

A final answer

upvoted 1 times

  **agrawalachin** 3 years, 7 months ago

A is correct

upvoted 1 times

A database specialist was alerted that a production Amazon RDS MariaDB instance with 100 GB of storage was out of space. In response, the database specialist modified the DB instance and added 50 GB of storage capacity. Three hours later, a new alert is generated due to a lack of free space on the same DB instance.

The database specialist decides to modify the instance immediately to increase its storage capacity by 20 GB.

What will happen when the modification is submitted?

- A. The request will fail because this storage capacity is too large.
- B. The request will succeed only if the primary instance is in active status.
- C. The request will succeed only if CPU utilization is less than 10%.
- D. The request will fail as the most recent modification was too soon.

Suggested Answer: B

Community vote distribution

D (100%)

🗳️ 👤 **Jaypdv** Highly Voted 👍 3 years, 7 months ago

D. answer. instance can't have any more storage modifications for six hours
upvoted 14 times

🗳️ 👤 **Satprave** Most Recent 🕒 2 years, 6 months ago

D - Storage shouldn't be extended immediately
upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: D

needs 6 hour gap
upvoted 4 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: D

Answer is D - cannot modify storage until EITHER 6 hours have passed OR the "storage-optimization" status is complete (instance will show "storage-optimization" happens after previous storage capacity has increased -it CAN take more than 6 hours)
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.StorageTypes.html
upvoted 4 times

🗳️ 👤 **kped21** 3 years, 3 months ago

D

Storage optimization can take several hours. You can't make further storage modifications for either six (6) hours or until storage optimization has completed on the instance, whichever is longer.
upvoted 2 times

🗳️ 👤 **GMartinelli** 3 years, 5 months ago

Selected Answer: D

Option D
upvoted 2 times

🗳️ 👤 **Scunningham99** 3 years, 6 months ago

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.StorageTypes.html - either 6 hours or when the first job completes - I would go with D too
upvoted 1 times

🗳️ 👤 **Aesthet** 3 years, 7 months ago

D

You can't make further storage modifications until six (6) hours after storage optimization has completed on the instance.
upvoted 2 times

🗳️ 👤 **agrawalachin** 3 years, 7 months ago

D is correct. 6 hours duration needs to pass


upvoted 1 times

  **shantest1** 3 years, 8 months ago

D. answer

I think it needs to pass 6 hours to increase another storage space increase.

upvoted 2 times

  **shantest1** 3 years, 7 months ago

Ignore the answer, that is for auto scaling, has to pass 6 hours.

upvoted 2 times

  **shantest1** 3 years, 7 months ago

Well, that condition applies both to manual as well as auto scaling, 6 hours has to pass. So I think it is D still

upvoted 1 times

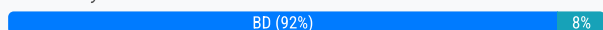
A company uses Amazon Aurora for secure financial transactions. The data must always be encrypted at rest and in transit to meet compliance requirements.

Which combination of actions should a database specialist take to meet these requirements? (Choose two.)

- A. Create an Aurora Replica with encryption enabled using AWS Key Management Service (AWS KMS). Then promote the replica to master.
- B. Use SSL/TLS to secure the in-transit connection between the financial application and the Aurora DB cluster.
- C. Modify the existing Aurora DB cluster and enable encryption using an AWS Key Management Service (AWS KMS) encryption key. Apply the changes immediately.
- D. Take a snapshot of the Aurora DB cluster and encrypt the snapshot using an AWS Key Management Service (AWS KMS) encryption key. Restore the snapshot to a new DB cluster and update the financial application database endpoints.
- E. Use AWS Key Management Service (AWS KMS) to secure the in-transit connection between the financial application and the Aurora DB cluster.

Suggested Answer: BC

Community vote distribution



shantest1 Highly Voted 3 years, 7 months ago

B and D.

upvoted 14 times

tugboat Highly Voted 3 years, 3 months ago

Selected Answer: BD

Per - <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Overview.Encryption.html>

Not A as - You can't create an encrypted Aurora Replica from an unencrypted Aurora DB cluster. You can't create an unencrypted Aurora Replica from an encrypted Aurora DB cluster.

B is good for in-transit replication

Not C as - You can't convert an unencrypted DB cluster to an encrypted one.

D as - You can restore an unencrypted snapshot to an encrypted Aurora DB cluster. To do this, specify a KMS key when you restore from the unencrypted snapshot.

Not E as - KMS does not perform encryption for data in transit or in motion. If you want to encrypt data while in transit, then you would need to use a different method such as SSL.

So, B and D is correct.

upvoted 6 times

redman50 Most Recent 2 years, 2 months ago

Selected Answer: AC

In Aurora you can encrypt at rest without copying the snapshot. So A and C for sure

upvoted 1 times

Piccaso 2 years, 1 month ago

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Overview.Encryption.html#Overview.Encryption.Limitations>

upvoted 1 times

dougporto1988 2 years, 2 months ago

Selected Answer: BD

I reckon is B and D

upvoted 1 times

lunt 3 years ago

Selected Answer: BD

B and D. D is right. Take snapshot of cluster > and (keyword here) > enable encryption. You cannot take a snapshot and encrypt it at the same time, this where the 'and' comes into play, you can encrypt just a snapshot + you can encrypt the snapshot on restore.

upvoted 2 times

🗨️ **novice_expert** 3 years ago

Selected Answer: BD

B. SSL/TLS is good for in-transit replication

D. as - You can restore an unencrypted snapshot to an encrypted Aurora DB cluster

upvoted 2 times

🗨️ **novice_expert** 3 years ago

D. as - You can NOT restore an unencrypted snapshot to an encrypted Aurora DB cluster

upvoted 1 times

🗨️ **kped21** 3 years, 3 months ago

B,D

C: Wrong, you cannot modify an unencrypted to encrypted

upvoted 3 times

🗨️ **awsmonster** 3 years, 4 months ago

A and B ..

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-replicas-adding.html>

D is incorrect:

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_CopySnapshot.html

For Amazon Aurora DB cluster snapshots, you can't encrypt an unencrypted DB cluster snapshot when you copy the snapshot.

upvoted 4 times

🗨️ **ChauPhan** 3 years, 6 months ago

B and D is correct

upvoted 1 times

🗨️ **Hits_23** 3 years, 7 months ago

B and D are correct choice.

upvoted 1 times

🗨️ **Jaypdv** 3 years, 8 months ago

BD

B. is obvious. For D. I thought it's possible to directly restore the unencrypted snapshot into an encrypted cluster so somehow one step looks unnecessary. But A, C and E are incorrect so I pick D. by default

upvoted 5 times

A company is running a website on Amazon EC2 instances deployed in multiple Availability Zones (AZs). The site performs a high number of repetitive reads and writes each second on an Amazon RDS for MySQL Multi-AZ DB instance with General Purpose SSD (gp2) storage. After comprehensive testing and analysis, a database specialist discovers that there is high read latency and high CPU utilization on the DB instance. Which approach should the database specialist take to resolve this issue without changing the application?



- A. Implement sharding to distribute the load to multiple RDS for MySQL databases.
- B. Use the same RDS for MySQL instance class with Provisioned IOPS (PIOPS) storage.
- C. Add an RDS for MySQL read replica.
- D. Modify the RDS for MySQL database class to a bigger size and implement Provisioned IOPS (PIOPS).

Suggested Answer: C

Community vote distribution

D (75%)

B (25%)

 **Scunningham99**  3 years, 6 months ago

I would go with D due to the high cpu utilisation; not C as this would require app to use read endpoint - therefore app change required
upvoted 15 times

 **redman50**  2 years, 2 months ago

Selected Answer: B

The high read latency and high CPU utilization on the Amazon RDS for MySQL Multi-AZ DB instance can be addressed by using Provisioned IOPS (PIOPS) storage with the same RDS for MySQL instance class. Therefore, option B is the correct approach to resolve the issue without changing the application. By using Provisioned IOPS (PIOPS), the database can deliver predictable and consistent I/O performance, which helps improve read latency and CPU utilization.

upvoted 1 times

 **synthia2** 2 years, 1 month ago

Multi-AZ is not for performance.

upvoted 1 times

 **saikirankshatriya** 2 years, 6 months ago

Selected Answer: D

B - doesnt incur downtime can change and improve performance but on the storage front

C - Needs to send high load queries to replicas

D - Chance of 60 second downtime due to change in instance class but performance improvement over CPU and storage.

upvoted 1 times

 **novice_expert** 3 years ago

Selected Answer: D

x C. Add an RDS for MySQL read replica. (Would need program to use Read End Point, else it is good solution)

D. Modify the RDS for MySQL database class to a bigger size and implement Provisioned IOPS (PIOPS). (because of high CPU + no change in code)

upvoted 2 times

 **RotterDam** 3 years, 2 months ago

CPU Utilization is High - This means the Database is underprovisioned - this drives the overall strategy to remediate the Instance class. (D) is the only option

upvoted 3 times

 **deepcloud** 3 years, 4 months ago

"without requiring program to be changed" - Adding read replica will need program to be changed to add replica endpoint. D should be the answer

upvoted 3 times

 **Amy2009** 3 years, 7 months ago

B should be the answer

upvoted 1 times

 **patricpotter1992** 3 years, 7 months ago

Correct Answer Should be D because of "Which approach should the database specialist to take to resolve this issue without changing the application?" WITHOUT CHANGING THE APPLICATION. Option C would need add the read-endpoints to the application.

upvoted 4 times

🗨️ 👤 **ChauPhan** 3 years, 7 months ago

After comprehensive testing and analysis, a database specialist discovers that there is high read latency and high CPU utilization on the DB instance.
=> C. Add read replicas for read SQL.

upvoted 1 times

🗨️ 👤 **sbhujbal** 3 years, 7 months ago

option c is out since "The site performs a high number of repetitive reads and writes". So mostly Option D is correct

upvoted 1 times

🗨️ 👤 **ChauPhan** 3 years, 7 months ago

After comprehensive testing and analysis, a database specialist discovers that there is high read latency and high CPU utilization on the DB instance

upvoted 2 times

🗨️ 👤 **AM** 3 years, 7 months ago

The only issue with C is that the application needs to be changed to point read traffic to read replica. This is not allowed as per the question. In that case the instance size increase will take care of the high CPU utilization. PIOPS will be bonus (As per the option D). I will go with option D.

upvoted 3 times

🗨️ 👤 **Aesthet** 3 years, 7 months ago

Both C and D will do the job.

"to resolve this issue without changing the application" - with C you have at least balance some or all reads across master and replica, add a connection string to read replica. So this makes D a correct answer (even though I prefer answer C over D if not that formulation).

upvoted 3 times

🗨️ 👤 **agrawalachin** 3 years, 7 months ago

C. The read latency and CPU util would reduce as read queries would be routed to Read replica

upvoted 1 times

🗨️ 👤 **manan728** 3 years, 7 months ago

It's debatable. BCD all look plausible options.

upvoted 1 times

🗨️ 👤 **manan728** 3 years, 7 months ago

Changing to C. There's no indication that the IO is high.

upvoted 1 times

🗨️ 👤 **notcloudguru** 3 years, 7 months ago

application needs to change if C, to read from replica

upvoted 1 times

🗨️ 👤 **shantest1** 3 years, 8 months ago

C. Read replica answer seems to be reasonable considering other options listed.

upvoted 1 times

A banking company recently launched an Amazon RDS for MySQL DB instance as part of a proof-of-concept project. A database specialist has configured automated database snapshots. As a part of routine testing, the database specialist noticed one day that the automated database snapshot was not created.

Which of the following are possible reasons why the snapshot was not created? (Choose two.)

- A. A copy of the RDS automated snapshot for this DB instance is in progress within the same AWS Region.
- B. A copy of the RDS automated snapshot for this DB instance is in progress in a different AWS Region.
- C. The RDS maintenance window is not configured.
- D. The RDS DB instance is in the STORAGE_FULL state.
- E. RDS event notifications have not been enabled.

Suggested Answer: AC

Community vote distribution

AD (93%)

7%

 **tugboat** Highly Voted 3 years, 3 months ago

Selected Answer: AD

Per - https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithAutomatedBackups.html

Your DB instance must be in the AVAILABLE state for automated backups to occur. Automated backups don't occur while your DB instance is in a state other than AVAILABLE, for example STORAGE_FULL.

Automated backups don't occur while a DB snapshot copy is running in the same AWS Region for the same DB instance.

upvoted 9 times

 **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: AD

"Automated backups follow these rules:

Your DB instance must be in the available state for automated backups to occur. Automated backups don't occur while your DB instance is in a state other than available, for example, storage_full.

Automated backups don't occur while a DB snapshot copy is running in the same AWS Region for the same database."

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithAutomatedBackups.html

upvoted 1 times

 **novice_expert** 3 years ago

Selected Answer: AD

A. Automated backups don't occur while a DB snapshot copy is running in the same AWS Region for the same DB instance.

D. DB instance must be in the AVAILABLE state


upvoted 2 times

 **user0001** 3 years, 2 months ago

Selected Answer: AC

since the database is not down, that is mean storage_full is not the case

upvoted 1 times

 **kped21** 3 years, 3 months ago

A,D

Automated backups follow these rules:

Your DB instance must be in the AVAILABLE state for automated backups to occur. Automated backups don't occur while your DB instance is in a state other than AVAILABLE, for example STORAGE_FULL.

Automated backups don't occur while a DB snapshot copy is running in the same AWS Region for the same DB instance.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithAutomatedBackups.html

upvoted 1 times

🗳️ 👤 **Imkjhgqysdfugshck** 3 years, 5 months ago

Selected Answer: AD

AD is the answer

upvoted 2 times

🗳️ 👤 **Suresh108** 3 years, 7 months ago

AD is the answer

upvoted 3 times

🗳️ 👤 **Aesthet** 3 years, 7 months ago

AD final answer

upvoted 1 times

🗳️ 👤 **manan728** 3 years, 7 months ago

AD is right.

upvoted 1 times

🗳️ 👤 **Jaypdy** 3 years, 8 months ago

Answer is AD

upvoted 4 times

🗳️ 👤 **[Removed]** 3 years, 8 months ago

I concur https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithAutomatedBackups.html

upvoted 4 times

🗳️ 👤 **gelsm** 3 years, 7 months ago

I agree with the link above

upvoted 1 times



An online shopping company has a large inflow of shopping requests daily. As a result, there is a consistent load on the company's Amazon RDS database. A database specialist needs to ensure the database is up and running at all times. The database specialist wants an automatic notification system for issues that may cause database downtime or for configuration changes made to the database. What should the database specialist do to achieve this? (Choose two.)

- A. Create an Amazon CloudWatch Events event to send a notification using Amazon SNS on every API call logged in AWS CloudTrail.
- B. Subscribe to an RDS event subscription and configure it to use an Amazon SNS topic to send notifications.
- C. Use Amazon SES to send notifications based on configured Amazon CloudWatch Events events.
- D. Configure Amazon CloudWatch alarms on various metrics, such as FreeStorageSpace for the RDS instance.
- E. Enable email notifications for AWS Trusted Advisor.

Suggested Answer: BC

Community vote distribution

BD (100%)

 **shantest1**  3 years, 7 months ago

B and D looks correct. With low storage instance will go out of service causing downtime.
Tricky at the same time easy question.
upvoted 18 times


 **Jaypdv**  3 years, 8 months ago

BD. Answer
upvoted 6 times

 **sachin**  2 years, 11 months ago

My take B and D
upvoted 1 times


 **novice_expert** 3 years ago

 BD


B. Subscribe to an RDS event subscription and configure it to use an Amazon SNS topic to send notifications.


D. Configure Amazon CloudWatch alarms on various metrics, such as FreeStorageSpace for the RDS instance.
upvoted 4 times

 **RotterDam** 3 years, 2 months ago

 BD

BD is correct
upvoted 1 times

 **tugboat** 3 years, 3 months ago

 BD

Valid options
upvoted 1 times

 **DuetRichard** 3 years, 7 months ago

A, D is answer
A -> configuration changes (Cloud trail can be used to detect api calls associated with any db change)
D -> cause database downtime (CloudWatch metrics)
upvoted 1 times

 **RotterDam** 3 years, 2 months ago

No for two reasons
1) Getting Cloudtrail for every event
2) You wont be alerted to Alarms on things like Low database storage space etc
upvoted 1 times

🗨️ 👤 **jove** 3 years, 6 months ago

A is sending notifications for every single API call logged in AWS CloudTrail. You don't want to do that :)
upvoted 6 times

🗨️ 👤 **Aesthet** 3 years, 7 months ago

DB final answer
upvoted 1 times

🗨️ 👤 **novak18** 3 years, 8 months ago

Shouldn't the answer be, A & B?

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/WhatIsCloudWatchEvents.html>

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_Events.html

upvoted 2 times

🗨️ 👤 **agrawalachin** 3 years, 7 months ago

A is incorrect. We don't need notification on every API call logged.
upvoted 5 times

A large company has a variety of Amazon DB clusters. Each of these clusters has various configurations that adhere to various requirements. Depending on the team and use case, these configurations can be organized into broader categories.

A database administrator wants to make the process of storing and modifying these parameters more systematic. The database administrator also wants to ensure that changes to individual categories of configurations are automatically applied to all instances when required.

Which AWS service or feature will help automate and achieve this objective?

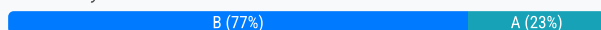
- A. AWS Systems Manager Parameter Store
- B. DB parameter group
- C. AWS Config
- D. AWS Secrets Manager

Suggested Answer: B

Reference:

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_WorkingWithParamGroups.html

Community vote distribution



Huy Highly Voted 3 years, 7 months ago

In AWS, DB Configurations are stored in DB Parameter Group. So B

upvoted 16 times

novak18 Highly Voted 3 years, 8 months ago

The answer should be C?

upvoted 6 times

tcl08 Most Recent 1 year, 3 months ago

There's Aurora Global Database allows a single Amazon Aurora database to span multiple AWS regions, enabling high performance of globally distributed applications.

upvoted 1 times

SamDDD 1 year, 10 months ago

Should be C - AWS Config: <https://aws.amazon.com/blogs/database/enforce-configuration-policies-for-your-amazon-rds-databases-using-aws-config/>

upvoted 2 times

sk1974 2 years, 3 months ago

<https://aws.amazon.com/blogs/database/enforce-configuration-policies-for-your-amazon-rds-databases-using-aws-config/> - I will go with C .Company has multiple DB instances.

upvoted 1 times

elf78 2 years, 10 months ago

Selected Answer: B

DB Parameter Group

upvoted 2 times

sachin 2 years, 11 months ago

My take is on A.

Certain paramters are only applied to certain instances ,, so different Parameter Group for different dbs.

On top DBA want to make sure certain configuration are applied commanly to all instance .

This can only be done though parameter store which can be implemented using Cloud Formation template.

Parameter Store, a capability of AWS Systems Manager, provides secure, hierarchical storage for configuration data management and secrets management. You can store data such as passwords, database strings, Amazon Machine Image (AMI) IDs, and license codes as parameter values. You can store values as plain text or encrypted data. You can reference Systems Manager parameters in your scripts, commands, SSM documents, and configuration and automation workflows by using the unique name that you specified when you created the parameter.

upvoted 1 times

db2luwdba 2 years, 10 months ago

Parameter hierarchies will ensure that configuration is properly categorized into wider groups and users.
The following example uses three hierarchy levels in the name to identify the following:

/Environment/Type of computer/Application/Data



/Dev/DBServer/MySQL/db-string13

upvoted 1 times

  **db2luwdba** 2 years, 10 months ago

So i believe Parameter store is correct.

upvoted 2 times

  **Dantas** 2 years, 12 months ago

Selected Answer: B

B, https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_WorkingWithParamGroups.html

upvoted 2 times

  **Mintwater** 2 years, 1 month ago

A DB parameter group is a collection of database engine parameter values that can be applied to one or more Amazon RDS, Amazon Aurora, or Amazon DocumentDB instances. By creating a DB parameter group, a database administrator can define and manage a set of database engine parameter values that apply to instances in a consistent manner.

With a DB parameter group, a database administrator can modify parameter values for multiple instances at once and have the changes automatically applied to all instances associated with the parameter group. This allows for centralized management and ensures that changes are applied consistently across all instances.

AWS Systems Manager Parameter Store, AWS Config, and AWS Secrets Manager are all AWS services that can be used to store and manage configuration parameters. However, these services are not specifically designed for managing database cluster configurations in a systematic manner and may not provide the same level of automation and consistency as a DB parameter group.

upvoted 2 times

  **novice_expert** 3 years ago

Selected Answer: B

Organizations create Standard Parameter Groups (for RDS) and Cluster/DB level Parameter Groups (for Aurora). These settings are applied uniformly across all instances implementing them.

upvoted 2 times

  **RotterDam** 3 years, 2 months ago

Got this question in my exam. (i cleared it). B is correct

upvoted 2 times

  **RotterDam** 3 years, 2 months ago

Selected Answer: B

Organizations create Standard Parameter Groups (for RDS) and Cluster/DB level Parameter Groups (for Aurora). These settings are applied uniformly across all instances implementing them.

Also For people choosing System Manager Parameter Store - the SSM does not integrate with RDS.

You can use Parameter Store parameters with other Systems Manager capabilities and AWS services to retrieve secrets and configuration data from a central store. Parameters work with Systems Manager capabilities such as Run Command, Automation, and State Manager, capabilities of AWS Systems Manager. You can also reference parameters in a number of other AWS services, including the following:

Amazon Elastic Compute Cloud (Amazon EC2)

Amazon Elastic Container Service (Amazon ECS)

AWS Secrets Manager

AWS Lambda

AWS CloudFormation

AWS CodeBuild

AWS CodePipeline

AWS CodeDeploy



upvoted 4 times

  **user0001** 3 years, 2 months ago

I vote for C since the requirement is to auto implement and validation.

B is good but it work at instance level and admin will need to manage them seperatly

upvoted 2 times



  **thelad** 3 years, 3 months ago

Selected Answer: A

<https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-parameter-store.html>

Who should use Parameter Store? - Any AWS customer who wants to have a centralized way to manage configuration data.

upvoted 2 times

  **jove** 3 years, 5 months ago

"... changes to are automatically implemented to all instances as necessary".. Can you use AWS Systems Manager Parameter Store for that? I don't think so. So, the correct answer should be B (DB parameter group)

upvoted 2 times

  **RotterDam** 3 years, 2 months ago

I agree- unless they are talking about CLOUDFORMATION and building Clusters and storing parameters centrally - this might make sense...I feel its PG.

upvoted 2 times

  **2025flakyt** 3 years, 5 months ago

B is the correct answer

Database parameters specify how the database is configured. For example, database parameters can specify the amount of resources, such as memory, to allocate to a database.

upvoted 1 times

  **GMartinelli** 3 years, 5 months ago

Selected Answer: A

Option A

upvoted 1 times

  **Sp230** 3 years, 6 months ago

A is the answer

upvoted 1 times

A company is developing a new web application. An AWS CloudFormation template was created as a part of the build process. Recently, a change was made to an AWS::RDS::DBInstance resource in the template. The CharacterSetName property was changed to allow the application to process international text. A change set was generated using the new template, which indicated that the existing DB instance should be replaced during an upgrade.

What should a database specialist do to prevent data loss during the stack upgrade?

- A. Create a snapshot of the DB instance. Modify the template to add the DBSnapshotIdentifier property with the ID of the DB snapshot. Update the stack.
- B. Modify the stack policy using the aws cloudformation update-stack command and the set-stack-policy command, then make the DB resource protected.
- C. Create a snapshot of the DB instance. Update the stack. Restore the database to a new instance.
- D. Deactivate any applications that are using the DB instance. Create a snapshot of the DB instance. Modify the template to add the DBSnapshotIdentifier property with the ID of the DB snapshot. Update the stack and reactivate the applications.

Suggested Answer: D

Reference:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-properties-rds-database-instance.html>

Community vote distribution

D (100%)

🗳️ 👤 **shantest1** Highly Voted 3 years, 7 months ago

D. appears to be correct based on the URL listed.

upvoted 6 times

🗳️ 👤 **Doox** Most Recent 1 year, 1 month ago

B :: via stack policy

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: D

Deactivate any applications that are using the DB instance.

-> Create a snapshot of the DB instance.

-> Modify the template to add the DBSnapshotIdentifier property with the ID of the DB snapshot.

-> Update the stack

-> and reactivate the applications.

upvoted 2 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

D seems to be the right answer

upvoted 1 times

🗳️ 👤 **TonyGe** 3 years, 7 months ago

D.

To preserve your data, perform the following procedure:

1. Deactivate any applications that are using the DB instance so that there's no activity on the DB instance.

2. Create a snapshot of the DB instance. For more information about creating DB snapshots

3. If you want to restore your instance using a DB snapshot, modify the updated template with your DB instance changes and add the DBSnapshotIdentifier property with the ID of the DB snapshot that you want to use

4. Update the stack.

upvoted 3 times

🗳️ 👤 **Hits_23** 3 years, 7 months ago

Correct answer is D

upvoted 1 times

🗳️ 👤 **Aesthet** 3 years, 7 months ago

D final answer
upvoted 2 times

A company recently acquired a new business. A database specialist must migrate an unencrypted 12 TB Amazon RDS for MySQL DB instance to a new AWS account. The database specialist needs to minimize the amount of time required to migrate the database. Which solution meets these requirements?

- A. Create a snapshot of the source DB instance in the source account. Share the snapshot with the destination account. In the target account, create a DB instance from the snapshot.
- B. Use AWS Resource Access Manager to share the source DB instance with the destination account. Create a DB instance in the destination account using the shared resource.
- C. Create a read replica of the DB instance. Give the destination account access to the read replica. In the destination account, create a snapshot of the shared read replica and provision a new RDS for MySQL DB instance.
- D. Use mysqldump to back up the source database. Create an RDS for MySQL DB instance in the destination account. Use the mysql command to restore the backup in the destination database.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **Zhongkai** Highly Voted 3 years, 7 months ago

Sharing an unencrypted manual DB snapshot enables authorized AWS accounts to directly restore a DB instance from the snapshot instead of taking a copy of it and restoring from that.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ShareSnapshot.html

However Resource Access Manager could not share non-Aurora cluster.

<https://docs.aws.amazon.com/ram/latest/userguide/shareable.html>

Hence I will go with A.

upvoted 10 times

🗳️ 👤 **sachin** Most Recent 2 years, 11 months ago

A is correct.

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: A

create snapshot -> share with destination account -> copy -> create instance from snapshot

upvoted 1 times

🗳️ 👤 **Dantas** 3 years, 2 months ago

Selected Answer: A

A -> The database professional must keep the migration time to a minimum.

upvoted 1 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: A

A but the answer choice is insufficient. With just this you will see a downtime. You need binlog replication for minimal downtime (or DMS)

upvoted 3 times

🗳️ 👤 **user0001** 3 years, 2 months ago

Selected Answer: A

answer a

upvoted 1 times

🗳️ 👤 **Shunpin** 3 years, 5 months ago

Ans: A

Mysqldump is too slow. If mysqlpump, I might consider this.

upvoted 3 times

🗳️ 👤 **toppic26** 3 years, 6 months ago

12 TB dump meaningless and slow, then you need to share? Answer is A.

upvoted 1 times

🗨️ 👤 **ChauPhan** 3 years, 7 months ago

A and D both are both but I think A is faster than D. mysqldump for 12T data is nearly impossible.

upvoted 3 times

🗨️ 👤 **Aesthet** 3 years, 7 months ago

A final answer

upvoted 3 times

🗨️ 👤 **AM** 3 years, 7 months ago

B looks good to me. Criteria is minimize migration time. 12TB snapshot will also take time.

upvoted 1 times

🗨️ 👤 **manan728** 3 years, 8 months ago

A is correct even though it is missing a crucial step of copying the snapshot in the target account. Simply sharing isn't sufficient.

upvoted 2 times

🗨️ 👤 **shantest1** 3 years, 8 months ago

A. snapshot is faster than the backup.

A appears to be the answer.

upvoted 4 times

A company has applications running on Amazon EC2 instances in a private subnet with no internet connectivity. The company deployed a new application that uses Amazon DynamoDB, but the application cannot connect to the DynamoDB tables. A developer already checked that all permissions are set correctly.

What should a database specialist do to resolve this issue while minimizing access to external resources?

- A. Add a route to an internet gateway in the subnet's route table.
- B. Add a route to a NAT gateway in the subnet's route table.
- C. Assign a new security group to the EC2 instances with an outbound rule to ports 80 and 443.
- D. Create a VPC endpoint for DynamoDB and add a route to the endpoint in the subnet's route table.

Suggested Answer: B

Community vote distribution

D (100%)

🗳️ 👤 **RotterDam** Highly Voted 3 years, 2 months ago

Got this question in my exam. (i cleared it). D is correct
upvoted 5 times

🗳️ 👤 **Zdujgr567783ff** 1 year, 11 months ago

how do u know this your answer was correct?
upvoted 2 times

🗳️ 👤 **jove** Highly Voted 3 years, 5 months ago

Selected Answer: D

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/vpc-endpoints-dynamodb.html>
upvoted 5 times

🗳️ 👤 **NishithShah** Most Recent 1 year, 5 months ago

D. Create a VPC endpoint for DynamoDB and add a route to the endpoint in the subnet's route table.
upvoted 1 times

🗳️ 👤 **Germaneli** 1 year, 8 months ago

Selected Answer: D

"What should a database specialist do to resolve this issue while minimizing access to external resources?"

x A. an internet gateway is actually maximising access to external resources - distractor.

x B. so is a NAT gateway.

x C. what do you want with ports 80 and 443? Distractor.

D. Create a VPC endpoint for DynamoDB: reasonable choice fulfilling the requirement.

upvoted 2 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: D

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/vpc-endpoints-dynamodb.html>

D. Create a VPC endpoint for DynamoDB -> add this endpoint in the subnet's route table.

upvoted 2 times

🗳️ 👤 **ChauPhan** 3 years, 7 months ago

D. Create a VPC endpoint for DynamoDB

upvoted 1 times

🗳️ 👤 **Aesthet** 3 years, 7 months ago

D final answer

upvoted 2 times

🗳️ 👤 **Aesthet** 3 years, 7 months ago



agree with D

upvoted 1 times

  **manan728** 3 years, 7 months ago

Yes D is right. You need a VPC endpoint in this case.

upvoted 2 times

  **Jaypdlv** 3 years, 7 months ago

D. Answer

upvoted 4 times

The Amazon CloudWatch metric for FreeLocalStorage on an Amazon Aurora MySQL DB instance shows that the amount of local storage is below 10 MB. A database engineer must increase the local storage available in the Aurora DB instance. How should the database engineer meet this requirement?

- A. Modify the DB instance to use an instance class that provides more local SSD storage.
- B. Modify the Aurora DB cluster to enable automatic volume resizing.
- C. Increase the local storage by upgrading the database engine version.
- D. Modify the DB instance and configure the required storage volume in the configuration section.

Suggested Answer: D

Community vote distribution

A (100%)

🗳️ 👤 **shantest1** Highly Voted 🍌 3 years, 8 months ago

A. answer

upvoted 7 times

🗳️ 👤 **shantest1** 3 years, 8 months ago

Local storage is the key here. Not the Database storage.

upvoted 4 times

🗳️ 👤 **Dantas** Most Recent 🕒 2 years, 12 months ago

Selected Answer: A

A, for local storage.

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: A

change to instance class that provides more local SSD storage.

<https://aws.amazon.com/premiumsupport/knowledge-center/aurora-mysql-local-storage/>

Local storage for each Aurora instance in the cluster, based on the instance class. This storage type and size is bound to the instance class, and can be changed only by moving to a larger DB instance class.

upvoted 3 times

🗳️ 👤 **kped21** 3 years, 3 months ago

A

You can increase the amount of free storage space for an instance by choosing a larger DB instance class for your instance.

upvoted 1 times

🗳️ 👤 **GMartinelli** 3 years, 5 months ago

Selected Answer: A

Option A

upvoted 2 times

🗳️ 👤 **Aesthet** 3 years, 7 months ago

A final answer

upvoted 2 times

🗳️ 👤 **Zhongkai** 3 years, 7 months ago

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.AuroraMySQL.Monitoring.Metrics.html> - A

upvoted 4 times

🗳️ 👤 **gelsm** 3 years, 6 months ago

"The amount of local storage available.

Unlike for other DB engines, for Aurora DB instances this metric reports the amount of storage available to each DB instance. This value depends

on the DB instance class (for pricing information, see the Amazon RDS product page). You can increase the amount of free storage space for an instance by choosing a larger DB instance class for your instance."

upvoted 2 times

  **manan728** 3 years, 8 months ago

Also I could remember a question on tier-0 vs tier-1 priority on Aurora, basically asking how to reduce failover time: if both master and replica are tier-0 or both master and replica should be tier-1.

upvoted 1 times

  **swarndeep** 3 years, 8 months ago

Answer A

upvoted 1 times

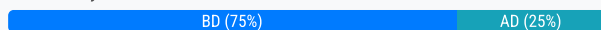
A company has an ecommerce web application with an Amazon RDS for MySQL DB instance. The marketing team has noticed some unexpected updates to the product and pricing information on the website, which is impacting sales targets. The marketing team wants a database specialist to audit future database activity to help identify how and when the changes are being made.

What should the database specialist do to meet these requirements? (Choose two.)

- A. Create an RDS event subscription to the audit event type.
- B. Enable auditing of CONNECT and QUERY_DML events.
- C. SSH to the DB instance and review the database logs.
- D. Publish the database logs to Amazon CloudWatch Logs.
- E. Enable Enhanced Monitoring on the DB instance.

Suggested Answer: AD

Community vote distribution



shantest1 Highly Voted 3 years, 8 months ago

B and D

upvoted 11 times

gelsm 3 years, 7 months ago

<https://aws.amazon.com/blogs/database/configuring-an-audit-log-to-capture-database-activities-for-amazon-rds-for-mysql-and-amazon-aurora-with-mysql-compatibility/>

upvoted 3 times

MultiAZ Most Recent 1 year, 4 months ago

Selected Answer: BD

B and D.

upvoted 1 times

aqiao 2 years, 1 month ago

Selected Answer: BD

A rules out because RDS event can not monitor data self change, it only monitor RDS config and parameter change, like DB snapshot, DB parameter group and so forth.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_Events.overview.html

upvoted 3 times

redman50 2 years, 2 months ago

Selected Answer: AD

CONNECT and QUERY_DML events are part of Aurora, not RDS, so answers A and D

upvoted 4 times

clarksu 2 years ago

CONNECT and QUERY_DML events are available for RDS. Your statement is false.

<https://aws.amazon.com/blogs/database/configuring-an-audit-log-to-capture-database-activities-for-amazon-rds-for-mysql-and-amazon-aurora-with-mysql-compatibility/>

upvoted 3 times

praffuln 3 years ago

Selected Answer: BD

enabling server_audit_events to auditing of CONNECT AND QUERY_DML events.

upvoted 1 times

novice_expert 3 years ago

Selected Answer: BD

x A. RDS event not helpful for DML

B. Enable auditing of CONNECT and QUERY_DML events.

x C. SSH not allowed

D. Publish the database logs to Amazon CloudWatch Logs.

x E. Enhanced Monitoring for performance

upvoted 3 times

🗨️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: BD

B and D is correct. You ENABLE audit logging via MARIADB_AUDIT_PLUGIN. You then modify the rds instance > select the option to export Audit Logs > this publishes it to Cloudwatch

upvoted 2 times

🗨️ 👤 **GMartinelli** 3 years, 5 months ago

Selected Answer: BD

Option B, D

upvoted 2 times

🗨️ 👤 **Aesthet** 3 years, 7 months ago

BD final answer

upvoted 2 times

🗨️ 👤 **agrawalachin** 3 years, 7 months ago

B & D are correct

upvoted 1 times

🗨️ 👤 **manan728** 3 years, 7 months ago

Why not A along with B?

upvoted 1 times

🗨️ 👤 **Justu** 3 years, 5 months ago

RDS Events are not for this, so right answer is B&D. A is Incorrect:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_Events.overview.html

upvoted 3 times

A large gaming company is creating a centralized solution to store player session state for multiple online games. The workload required key-value storage with low latency and will be an equal mix of reads and writes. Data should be written into the AWS Region closest to the user across the games' geographically distributed user base. The architecture should minimize the amount of overhead required to manage the replication of data between Regions.

Which solution meets these requirements?

- A. Amazon RDS for MySQL with multi-Region read replicas
- B. Amazon Aurora global database
- C. Amazon RDS for Oracle with GoldenGate
- D. Amazon DynamoDB global tables

Suggested Answer: A

Community vote distribution

D (100%)

 **shantest1** Highly Voted 3 years, 7 months ago

D. Answer

Key Value pair - DynamoDB Global tables

upvoted 14 times

 **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: D

D. Amazon DynamoDB global tables


key-value NoSQL, multi-region, low latency.

upvoted 1 times

 **Pranava_GCP** 1 year, 8 months ago


<https://aws.amazon.com/dynamodb/global-tables/>

upvoted 1 times

 **sriexam** 2 years, 10 months ago


I see the answer is showing as A. How is it justified for Key value Pair. Answer would be D. Please provide reason for answers .. also that would be helpful

upvoted 1 times

 **im_not_robot** 2 years, 3 months ago

you can not trust provided answer on this site

upvoted 4 times

 **eji** 2 years, 10 months ago


There is no Aurora global database, so the answer is D

upvoted 1 times

 **shammous** 2 years, 9 months ago

There is

upvoted 3 times

 **niau** 2 years, 11 months ago

Selected Answer: D

D. DynamoDB for sure

upvoted 1 times

 **novice_expert** 3 years ago

Selected Answer: D

Dynamo Global

https://aws.amazon.com/dynamodb/?nc1=h_ls

upvoted 1 times

🗲️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: D

DynamoDB Global tables is obvious for Key Value pair
upvoted 2 times

🗲️ 👤 **kped21** 3 years, 3 months ago

D, low latency and key-value
upvoted 3 times

🗲️ 👤 **Imkjhgqysdfugshck** 3 years, 5 months ago

Selected Answer: D

D. Answer
upvoted 3 times

🗲️ 👤 **akiraklaus** 3 years, 5 months ago

https://aws.amazon.com/dynamodb/?nc1=h_ls
D
upvoted 1 times

🗲️ 👤 **topic26** 3 years, 6 months ago

D. Obviously not MySQL: Question: low-latency key-value storage. This is DynamoDB local writes.
upvoted 1 times

🗲️ 👤 **Aesthet** 3 years, 6 months ago

D final answer
upvoted 2 times

🗲️ 👤 **manan728** 3 years, 7 months ago

Prepare on Aurora upgrades both minor and major versions. There were a few questions on the topic today at my actual test.
upvoted 2 times

🗲️ 👤 **Jaypdv** 3 years, 8 months ago

Answer is D.
B. could work, but the question specifically asks for a key-value store, hence DynamoDB
upvoted 3 times

🗲️ 👤 **Jiang_aws1** 2 years, 7 months ago

B not work due to "Aurora global database" Replic DB is ready only
upvoted 2 times

🗲️ 👤 **novak18** 3 years, 8 months ago

Since the question mentions key-value and also about gaming, I think the answer should be D in that case
upvoted 1 times

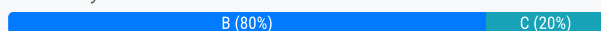
A company is running an on-premises application comprised of a web tier, an application tier, and a MySQL database tier. The database is used primarily during business hours with random activity peaks throughout the day. A database specialist needs to improve the availability and reduce the cost of the MySQL database tier as part of the company's migration to AWS.

Which MySQL database option would meet these requirements?

- A. Amazon RDS for MySQL with Multi-AZ
- B. Amazon Aurora Serverless MySQL cluster
- C. Amazon Aurora MySQL cluster
- D. Amazon RDS for MySQL with read replica

Suggested Answer: C

Community vote distribution



shantest1 Highly Voted 3 years, 8 months ago

B. Aurora Serverless

two requirements, reduce the cost and availability. Aurora Serverless is cheaper compared to Aurora cluster
upvoted 13 times

guru_ji 3 years, 7 months ago

Why not A ?

Amazon RDS for MySQL with Multi-AZ

upvoted 1 times

GMartinelli 3 years, 6 months ago

The question asks for minimal costs. Multi-AZ is super expensive, literally double the price of single AZ

upvoted 2 times

guru_ji 3 years, 7 months ago

ok.. B looks good.

primarily during business hours with random activity peaks throughout the day

upvoted 1 times

user0001 3 years, 2 months ago

A only provide HA ad not cost effective

upvoted 3 times

2025flakyt Highly Voted 3 years, 5 months ago

B is correct

Amazon Aurora Serverless v1 is a simple, cost-effective option for infrequent, intermittent, or unpredictable workloads.

<https://aws.amazon.com/rds/aurora/serverless/>

upvoted 6 times

Pranava_GCP Most Recent 1 year, 8 months ago

Selected Answer: B

B. Aurora Serverless

"database is used primarily during business hours with random activity peaks throughout the day" is key to select aurora serverless.

upvoted 1 times

Pankaj24hrs 2 years ago

B. Aurora Serverless

upvoted 1 times

praffuln 2 years, 4 months ago

Selected Answer: B

"database is used primarily during business hours with random activity peaks throughout the day" is key to select aurora serverless.
upvoted 3 times

  **Jiang_aws1** 2 years, 8 months ago



Selected Answer: C

"MySQL database tier as part of the company's migration to AWS" so we need to keep "tier" for DB-host so we can't use Serverless . "C" is correct
upvoted 2 times

  **novice_expert** 3 years ago



Selected Answer: B

accessed during business hours, with occasional bursts of activity throughout the day
increase the availability and minimize the cost
upvoted 2 times


  **tugboat** 3 years, 3 months ago

Selected Answer: B

B for occasional queries and minimal cost
upvoted 2 times

  **yahoos** 3 years, 3 months ago



B
Aurora Serverless MySQL cluster
upvoted 1 times

  **Nots** 3 years, 7 months ago

Considering the use cases described below, it is assumed that there is little traffic in the case of production use.
Also, since V2 is currently in preview, it cannot be expected to be in production.
Considering these, the Aurora DB cluster is considered to be a candidate, although the cost is higher than that of serverless.
<https://aws.amazon.com/rds/aurora/serverless/>
upvoted 1 times

  **ChauPhan** 3 years, 7 months ago

B. Aurora Serverless
upvoted 1 times

  **Aesthet** 3 years, 7 months ago

B final answer
upvoted 2 times

A company wants to migrate its Microsoft SQL Server Enterprise Edition database instance from on-premises to AWS. A deep review is performed and the AWS Schema Conversion Tool (AWS SCT) provides options for running this workload on Amazon RDS for SQL Server Enterprise Edition, Amazon RDS for SQL Server Standard Edition, Amazon Aurora MySQL, and Amazon Aurora PostgreSQL. The company does not want to use its own SQL server license and does not want to change from Microsoft SQL Server.

What is the MOST cost-effective and operationally efficient solution?

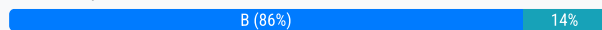
- A. Run SQL Server Enterprise Edition on Amazon EC2.
- B. Run SQL Server Standard Edition on Amazon RDS.
- C. Run SQL Server Enterprise Edition on Amazon RDS.
- D. Run Amazon Aurora MySQL leveraging SQL Server on Linux compatibility libraries.

Suggested Answer: D

Reference:

https://docs.aws.amazon.com/SchemaConversionTool/latest/userguide/CHAP_Welcome.html

Community vote distribution



Jaypdv Highly Voted 3 years, 7 months ago

B.

SCT assessment says that you can use MSSQL standard edition. Since they want to stay on SQL Server, that's the cheapest option and migrating is operationally easy

upvoted 14 times

shantest1 3 years, 7 months ago

re-read that part, Standard Edition looks suffice.

B. Answer.

upvoted 2 times

Pranava_GCP Most Recent 1 year, 9 months ago

Selected Answer: B

B. Run SQL Server Standard Edition on Amazon RDS.

upvoted 1 times

Bota5ky 1 year, 10 months ago

Selected Answer: C

They don't want to change version. And rds provides the license.

upvoted 1 times

Zdujgfr567783ff 1 year, 11 months ago

Selected Answer: B

Standard edition on RDS

upvoted 1 times

novice_expert 3 years ago

Selected Answer: B

B. Run SQL Server Standard Edition on Amazon RDS.

upvoted 1 times

jove 3 years, 5 months ago

Selected Answer: B

OPERATIONALLY EFFECTIVE = SQL Server on RDS

MOST COST-EFFECTIVE = Standard Edition

upvoted 3 times

grekh001 3 years, 6 months ago

This link seems to indicate that more information is required to determine if the Enterprise instance is a candidate for downgrading to Standard.
<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/determine-whether-your-microsoft-sql-server-database-can-be-downgraded-from-enterprise-to-standard-edition.html>

upvoted 1 times

🗨️ 👤 **sbhujbal** 3 years, 6 months ago

When Enterprise license is included in RDS price at FREE why to go to standard edition and loose all great features so Option C is correct.

upvoted 1 times

🗨️ 👤 **ExtHo** 3 years, 6 months ago

Try to calculate price Enterprise will cost you almost double for same specification

<https://calculator.aws/#/createCalculator/RDSSQLServer>

upvoted 4 times

🗨️ 👤 **Aesthet** 3 years, 7 months ago

B final answer

upvoted 3 times

🗨️ 👤 **Zhongkai** 3 years, 7 months ago

I will go with C as Standard Edition is not operationally efficient as Enterprise Edition. Check <https://docs.microsoft.com/en-us/sql/sql-server/editions-and-components-of-sql-server-2017?view=sql-server-ver15>

upvoted 2 times

🗨️ 👤 **shantest1** 3 years, 7 months ago

C. Run SQL Server Enterprise Edition on AWS RDS

License is included with the RDS price

Need Enterprise to Enterprise edition matching.

On EC2 you will use your own license.

upvoted 3 times

🗨️ 👤 **Suresh108** 3 years, 7 months ago

if standard edition is sufficient for the DB why to go for Enterprise edition and pay more??

i think B is right

upvoted 2 times

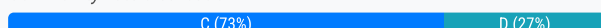
A company's ecommerce website uses Amazon DynamoDB for purchase orders. Each order is made up of a Customer ID and an Order ID. The DynamoDB table uses the Customer ID as the partition key and the Order ID as the sort key.

To meet a new requirement, the company also wants the ability to query the table by using a third attribute named Invoice ID. Queries using the Invoice ID must be strongly consistent. A database specialist must provide this capability with optimal performance and minimal overhead. What should the database administrator do to meet these requirements?

- A. Add a global secondary index on Invoice ID to the existing table.
- B. Add a local secondary index on Invoice ID to the existing table.
- C. Recreate the table by using the latest snapshot while adding a local secondary index on Invoice ID.
- D. Use the partition key and a FilterExpression parameter with a filter on Invoice ID for all queries.

Suggested Answer: C

Community vote distribution



🗳️ **shantest1** Highly Voted 3 years, 7 months ago

C. Answer
upvoted 10 times

🗳️ **jove** Highly Voted 3 years, 5 months ago

Selected Answer: C

Answer is C ;
- It has to be a local secondary index
- Secondary indexes can only be created when the table is created.
upvoted 5 times

🗳️ **jitesh_k** Most Recent 1 year, 5 months ago

What is wrong with D?
upvoted 1 times

🗳️ **Pranava_GCP** 1 year, 8 months ago

Selected Answer: C

C. Recreate table while adding a local secondary index

- strongly consistent has to be local secondary index
- adding a local secondary index requires recreating of the table
upvoted 1 times

🗳️ **Pankaj24hrs** 2 years ago

C
Due to clause "minimal overhead".

if you obtain 100KB of data in step 1 and filter it down to 10KB in step 2, you'll use the read capacity units for 100KB of data instead of the 10KB that was filtered down. Further, a 1MB limit is applied to all operations, regardless of the read capacity units on a table.
upvoted 1 times

🗳️ **SeemaDataReader** 2 years, 1 month ago

Selected Answer: D

A. Add a global secondary index on Invoice ID to the existing table. - GSI doesn't support strong consistency
B. Add a local secondary index on Invoice ID to the existing table. - LSI can be added only while creating the table. LSI Partition key has to be same as the main table's partition key i.e customer id
C. Recreate the table by using the latest snapshot while adding a local secondary index on Invoice ID. - LSI Partition key has to be same as the main table's partition key i.e customer id
D. Use the partition key and a FilterExpression parameter with a filter on Invoice ID for all queries. filter queries can be strongly consistent. If you require strongly consistent reads, set the ConsistentRead parameter to true in the Query request.
<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Query.html>

upvoted 4 times

🗳️ 👤 **yyy** 2 years ago

Ans is D.

You are right, C is wrong

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: C

Queries that make use of the Invoice ID must be very consistent => recreate table, as Local secondary index can only be created while creating the Dynamodb table.

global secondary indexes are eventual consistent only

upvoted 2 times

🗳️ 👤 **GMartinelli** 3 years, 5 months ago

Selected Answer: C

Option C

upvoted 3 times

🗳️ 👤 **ChauPhan** 3 years, 7 months ago

Why not A?. <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GSI.html>

upvoted 2 times

🗳️ 👤 **faramawi** 3 years, 6 months ago

According to below link "Strongly consistent reads are not supported on global secondary indexes"

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

upvoted 2 times

🗳️ 👤 **Hits_23** 3 years, 7 months ago

C is correct, as Local secondary index can only be created while creating the Dynamodb table. and query needs to use third attribute on top of primary and sort key, so Local Secondary index has primary and sort key as well as the third attribute. Global secondary index can be created without primary and sort key

upvoted 1 times

🗳️ 👤 **Suresh108** 3 years, 7 months ago

cccccc

upvoted 1 times

🗳️ 👤 **Aesthet** 3 years, 7 months ago

C final answer

upvoted 2 times

🗳️ 👤 **db_interest** 3 years, 7 months ago

C is correct. Local indexes provide strong consistency and cannot be created on existing tables.

upvoted 3 times

A company wants to migrate its on-premises MySQL databases to Amazon RDS for MySQL. To comply with the company's security policy, all databases must be encrypted at rest. RDS DB instance snapshots must also be shared across various accounts to provision testing and staging environments.

Which solution meets these requirements?

- A. Create an RDS for MySQL DB instance with an AWS Key Management Service (AWS KMS) customer managed CMK. Update the key policy to include the Amazon Resource Name (ARN) of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.
- B. Create an RDS for MySQL DB instance with an AWS managed CMK. Create a new key policy to include the Amazon Resource Name (ARN) of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.
- C. Create an RDS for MySQL DB instance with an AWS owned CMK. Create a new key policy to include the administrator user name of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.
- D. Create an RDS for MySQL DB instance with an AWS CloudHSM key. Update the key policy to include the Amazon Resource Name (ARN) of the other AWS accounts as a principal, and then allow the kms:CreateGrant action.

Suggested Answer: A

Reference:

<https://docs.aws.amazon.com/kms/latest/developerguide/grants.html>

Community vote distribution

A (100%)

🗳️ **Zhongkai** Highly Voted 3 years, 8 months ago

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ShareSnapshot.html - A

upvoted 5 times

🗳️ **Hisayuki** Most Recent 1 year, 4 months ago

Selected Answer: A

- KMS Customer managed CMK: The key is stored in your account that you created, own, and manage.
- KMS AWS managed CMK: The key is stored in your account and is managed by AWS Key Management Service.

So you should choose Customer managed CMK

upvoted 1 times

🗳️ **novice_expert** 3 years ago

Selected Answer: A

- Need customer managed CMK for sharing (B and C are out), C is also out for administrator user
- AWS CloudHSM key is also AWS generated (D is out)
- policy update to include ARN of other AWS accounts as principal
- CreateGrant action

upvoted 4 times

🗳️ **Dantas** 3 years, 2 months ago

Selected Answer: A

A!

You can't share a snapshot that has been encrypted using the default KMS key of the AWS account that shared the snapshot, therefore it must be encrypted with a customer managed key. https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ShareSnapshot.html

upvoted 4 times

🗳️ **RotterDam** 3 years, 2 months ago

Got this question in my exam. (i cleared it). A is correct

upvoted 1 times

🗳️ **kped21** 3 years, 3 months ago

A: To allow another AWS account access to a KMS key, update the key policy for the KMS key. You update it with the Amazon Resource Name (ARN) of the AWS account that you are sharing to as Principal in the KMS key policy. Then you allow the kms:CreateGrant action.

upvoted 2 times

🗨️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: A

KMS managed CMK

upvoted 1 times

🗨️ 👤 **jove** 3 years, 6 months ago

Selected Answer: A

CMK is needed for data share and you just need to update the key policy.

upvoted 2 times

🗨️ 👤 **ChauPhan** 3 years, 7 months ago

Agree with A

upvoted 1 times

🗨️ 👤 **Aesthet** 3 years, 7 months ago

A final answer

upvoted 2 times

🗨️ 👤 **shantest1** 3 years, 8 months ago

A. Answer

Key to the answer CMK - Customer managed Key - If I am not wrong

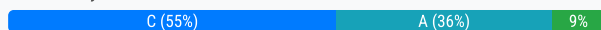
upvoted 3 times

A retail company manages a web application that stores data in an Amazon DynamoDB table. The company is undergoing account consolidation efforts. A database engineer needs to migrate the DynamoDB table from the current AWS account to a new AWS account. Which strategy meets these requirements with the LEAST amount of administrative work?

- A. Use AWS Glue to crawl the data in the DynamoDB table. Create a job using an available blueprint to export the data to Amazon S3. Import the data from the S3 file to a DynamoDB table in the new account.
- B. Create an AWS Lambda function to scan the items of the DynamoDB table in the current account and write to a file in Amazon S3. Create another Lambda function to read the S3 file and restore the items of a DynamoDB table in the new account.
- C. Use AWS Data Pipeline in the current account to export the data from the DynamoDB table to a file in Amazon S3. Use Data Pipeline to import the data from the S3 file to a DynamoDB table in the new account.
- D. Configure Amazon DynamoDB Streams for the DynamoDB table in the current account. Create an AWS Lambda function to read from the stream and write to a file in Amazon S3. Create another Lambda function to read the S3 file and restore the items to a DynamoDB table in the new account.

Suggested Answer: C

Community vote distribution



novak18 Highly Voted 3 years, 8 months ago

I think the answer is C

<https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-cross-account-migration/>

upvoted 11 times

DevoteamAnalytix 2 years, 10 months ago

For me it is A because it seems to be easier with Glue than Data Pipeline ("with the MINIMUM amount of administrative work")

GLUE: <https://aws.amazon.com/de/premiumsupport/knowledge-center/dynamodb-cross-account-migration/>

DATA PIPELINE: <https://aws.amazon.com/de/premiumsupport/knowledge-center/data-pipeline-account-access-dynamodb-s3/>

upvoted 4 times

roymunson Most Recent 1 year, 7 months ago

Selected Answer: C

Create a DynamoDB table in your source account.

Create an Amazon Simple Storage Service (Amazon S3) bucket in the destination account.

Attach an AWS Identity and Access Management (IAM) policy to the Data Pipeline default roles in the source account.

Create an S3 bucket policy in the destination account.

Create and activate a pipeline in the source account.

Create a DynamoDB table in the destination account.

Restore the DynamoDB export in the destination account.

<https://aws.amazon.com/de/blogs/database/how-to-migrate-amazon-dynamodb-tables-from-one-aws-account-to-another-with-aws-data-pipeline/>

upvoted 2 times

alexpl 1 year, 7 months ago

Selected Answer: C

The correct answer is C, because:

<https://aws.amazon.com/blogs/database/how-to-migrate-amazon-dynamodb-tables-from-one-aws-account-to-another-with-aws-data-pipeline/>

upvoted 1 times

thuyeinaung 1 year, 8 months ago

Selected Answer: A

I think it is A

upvoted 1 times

Germaneli 1 year, 8 months ago

Selected Answer: A

Migrate your DynamoDB table to a different AWS account with one of these methods that suit your use case:
<https://repost.aws/knowledge-center/dynamodb-cross-account-migration>

- AWS Backup
 - DynamoDB import and export to Amazon Simple Storage Service (Amazon S3)
 - Amazon S3 and AWS Glue
 - Amazon EMR
- upvoted 1 times

🗨️ **Monknill** 1 year, 10 months ago

C looks like the best option

<https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-cross-account-migration/>
upvoted 1 times

🗨️ **milan9527** 1 year, 11 months ago

Selected Answer: A

A. Why not?

upvoted 2 times

🗨️ **clarksu** 2 years ago

Selected Answer: C

You can migrate your DynamoDB tables to a different AWS account by choosing one of the following methods depending on your use case:

- AWS Backup
 - DynamoDB import and export to Amazon Simple Storage Service (Amazon S3)
 - Amazon S3 and AWS Glue
 - AWS Data Pipeline
 - Amazon EMR
- upvoted 1 times

🗨️ **Paulv82003** 1 year, 12 months ago

Amazon S3 and AWS Glue is on your list, listed before AWS Data Pipeline, and yet you select C?

upvoted 2 times

🗨️ **redman50** 2 years, 2 months ago

Selected Answer: D

Configure Amazon DynamoDB Streams for the DynamoDB table in the current account. Create an AWS Lambda function to read from the stream and write to a file in Amazon S3. Create another Lambda function to read the S3 file and restore the items to a DynamoDB table in the new account. This approach leverages DynamoDB Streams, which captures item-level modifications to a table, including creates, updates, and deletes. The DynamoDB Streams data can be used to replicate data in near real-time across different AWS accounts. The approach allows for minimal administrative work as it only requires the creation of two Lambda functions to read and write to S3 and DynamoDB tables, respectively. This approach also ensures that data consistency is maintained during the migration process.

upvoted 1 times

🗨️ **sk1974** 2 years, 3 months ago

<https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-cross-account-migration/> . I initially thought answer was 'A' since answer had the word 'existing blueprint' . but , I went for C based on the link pasted above. Scroll to the 'Data Pipeline' section in there .

upvoted 1 times

🗨️ **jjyy80** 2 years, 5 months ago

DATAPIEPLINE is the correct one "Note: The destination account can't access the DynamoDB data in S3 bucket. To work with the data, restore it to a DynamoDB table. Data Pipeline provides the easiest method to move the table with the least manual effort. However, there are fewer options for customization."

upvoted 2 times

🗨️ **novice_expert** 3 years ago

Selected Answer: C

<https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-cross-account-migration/>

export dynamoDB to S3 in other account -> use Glue job (or data pipeline or EMR) to import data

C. Use AWS Data Pipeline in the current account to export the data from the DynamoDB table to a file in Amazon S3. Use Data Pipeline to import the data from the S3 file to a DynamoDB table in the new account.

upvoted 2 times

🗨️ 👤 **johnconnor** 3 years, 6 months ago

Why not A? being glue serverless wouldn't it be easier to do this way?

upvoted 2 times

🗨️ 👤 **ChauPhan** 3 years, 6 months ago

Agree with C, LEAST amount of work.

upvoted 2 times

🗨️ 👤 **AM** 3 years, 7 months ago

I agree with C and also that the question is a bit ambiguous

upvoted 1 times

🗨️ 👤 **Aesthet** 3 years, 7 months ago

C

<https://aws.amazon.com/premiumsupport/knowledge-center/data-pipeline-account-access-dynamodb-s3/>

upvoted 4 times

🗨️ 👤 **manan728** 3 years, 7 months ago

A seems to be the answer.

<https://docs.aws.amazon.com/glue/latest/dg/aws-glue-programming-etl-dynamo-db-cross-account.html>

upvoted 2 times

A company uses the Amazon DynamoDB table contractDB in us-east-1 for its contract system with the following schema: orderID (primary key) timestamp (sort key) contract (map) createdBy (string) customerEmail (string)

After a problem in production, the operations team has asked a database specialist to provide an IAM policy to read items from the database to debug the application. In addition, the developer is not allowed to access the value of the customerEmail field to stay compliant.

Which IAM policy should the database specialist use to achieve these requirements?

A.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Allow",
      "Action": [
        "dynamodb: Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "orderID"
            "timestamp",
            "contract",
            "createdBy"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```

B.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Allow",
      "Action": [
        "dynamodb: Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "customerEmail"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}

```

C.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Deny",
      "Action": [
        "dynamodb: Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "customerEmail"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}

```


```
D.
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Deny",
      "Action": [
        "dynamodb: Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "orderId",
            "timestamp",
            "contract",
            "createdBy"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```

Suggested Answer: A

 **Jaypdv** Highly Voted 3 years, 7 months ago

A. Answer

upvoted 9 times

 **SachinGoel** Most Recent 2 years, 4 months ago

A is correct as per IAM guideline

upvoted 1 times

 **novice_expert** 3 years ago

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/specifying-conditions.html>


You use the IAM Condition element to implement a fine-grained access control policy.

upvoted 3 times

 **learnazureportal** 3 years, 6 months ago

The correct answer is A. you have access to all columns except CustomerEmail!

upvoted 2 times

 **Aesthet** 3 years, 6 months ago

A final answer

upvoted 2 times

 **Huy** 3 years, 7 months ago

A.

C would be correct if using `ForAnyValue` instead of `ForAllValues`
upvoted 3 times

A company has an application that uses an Amazon DynamoDB table to store user data. Every morning, a single-threaded process calls the DynamoDB API Scan operation to scan the entire table and generate a critical start-of-day report for management. A successful marketing campaign recently doubled the number of items in the table, and now the process takes too long to run and the report is not generated in time. A database specialist needs to improve the performance of the process. The database specialist notes that, when the process is running, 15% of the table's provisioned read capacity units (RCUs) are being used.

What should the database specialist do?

- A. Enable auto scaling for the DynamoDB table.
- B. Use four threads and parallel DynamoDB API Scan operations.
- C. Double the table's provisioned RCUs.
- D. Set the Limit and Offset parameters before every call to the API.

Suggested Answer: B

Community vote distribution

B (100%)

  **Jaypdr**  3 years, 7 months ago

B.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Scan.html#Scan.ParallelScan>

upvoted 13 times

  **MultiAZ**  1 year, 4 months ago

Selected Answer: B

Answer is B

upvoted 1 times

  **Pranava_GCP** 1 year, 8 months ago

Selected Answer: B

B. Use four threads and parallel DynamoDB API Scan operations.

upvoted 1 times

  **SachinGoel** 2 years, 4 months ago

Selected Answer: B

More thread is a key

upvoted 1 times

  **db2luwdba** 2 years, 10 months ago

B make sense but this has to done programmatically by the application team to process read by multi threaded application calls

upvoted 1 times

  **novice_expert** 3 years ago

Selected Answer: B



15% means max 6 threads possible, 4 are good

upvoted 2 times

  **ChauPhan** 3 years, 6 months ago

Every morning, a single-threaded process calls the DynamoDB API Scan operation to scan the entire table ==> B

upvoted 1 times

  **Aesthet** 3 years, 7 months ago

B final answer


upvoted 2 times

  **novak18** 3 years, 7 months ago

<https://aws.amazon.com/blogs/database/amazon-dynamodb-auto-scaling-performance-and-cost-optimization-at-any-scale/#:~:text=To%20configure%20auto%20scaling%20in,alarms%20that%20track%20consumed%20capacity.>



Answer should be A?

upvoted 1 times

  **awsmonster** 3 years, 4 months ago

15% utilization of the allocated RCUs does not required to scale the Dynamo DB.

upvoted 2 times

  **jove** 3 years, 5 months ago

No, it is B

upvoted 1 times

A company is building a software as a service application. As part of the new user sign-on workflow, a Python script invokes the CreateTable operation using the Amazon DynamoDB API. After the call returns, the script attempts to call PutItem. Occasionally, the PutItem request fails with a ResourceNotFoundException error, which causes the workflow to fail. The development team has confirmed that the same table name is used in the two API calls. How should a database specialist fix this issue?

- A. Add an allow statement for the dynamodb:PutItem action in a policy attached to the role used by the application creating the table.
- B. Set the StreamEnabled property of the StreamSpecification parameter to true, then call PutItem.
- C. Change the application to call DescribeTable periodically until the TableStatus is ACTIVE, then call PutItem.
- D. Add a ConditionExpression parameter in the PutItem request.

Suggested Answer: D

Reference:

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

Community vote distribution

C (83%)

D (17%)

🗳️ **Jaypdr** Highly Voted 3 years, 7 months ago

C. Answer
upvoted 11 times

🗳️ **MultiAZ** Most Recent 1 year, 4 months ago

Selected Answer: C

Answer C. The error happens OCCASIONALLY. So you need to wait for the table creation to complete.
upvoted 1 times

🗳️ **Pranava_GCP** 1 year, 8 months ago

Selected Answer: C

C. Change the application to call DescribeTable periodically until the TableStatus is ACTIVE, then call PutItem.

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

"ResourceNotFoundException

The operation tried to access a nonexistent table or index. The resource might not be specified correctly, or its status might not be ACTIVE."
upvoted 1 times

🗳️ **SamDDD** 1 year, 10 months ago

Selected Answer: C

DynamoDB responds with this error when you're trying to run an operation against non-existent or not active table.
upvoted 1 times

🗳️ **santosrk** 1 year, 11 months ago

Selected Answer: D

ConditionExpression

A condition that must be satisfied in order for a conditional PutItem operation to succeed.

An expression can contain any of the following:

Functions: attribute_exists | attribute_not_exists | attribute_type | contains | begins_with | size

These function names are case-sensitive.

Comparison operators: = | <> | < | > | <= | >= | BETWEEN | IN

Logical operators: AND | OR | NOT

upvoted 1 times

🗳️ 👤 **santosrk** 1 year, 11 months ago

D. answer. ConditionExpression

A condition that must be satisfied in order for a conditional PutItem operation to succeed.

An expression can contain any of the following:

Functions: attribute_exists | attribute_not_exists | attribute_type | contains | begins_with | size

These function names are case-sensitive.

Comparison operators: = | <> | < | > | <= | >= | BETWEEN | IN

Logical operators: AND | OR | NOT

upvoted 1 times

🗳️ 👤 **Kodoma** 1 year, 11 months ago

Selected Answer: C

The answer is C.

upvoted 1 times

🗳️ 👤 **TL12345** 2 years, 7 months ago

Answer is C.

ResourceNotFoundException

The operation tried to access a nonexistent table or index. The resource might not be specified correctly, or its status might not be ACTIVE.

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

upvoted 1 times

🗳️ 👤 **sachin** 2 years, 10 months ago

Condition Expression is used for PutItem DML conditions. you can specify a condition expression to determine which items should be modified. If the condition expression evaluates to true, the operation succeeds; otherwise, the operation fails.

To evaluate if the update or new entry should be made if the key attributes are same.

example : The PutItem operation overwrites an item with the same key (if it exists). If you want to avoid this, use a condition expression. This allows the write to proceed only if the item in question does not already have the same key.

C is correct.

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: C

TableStatus should be ACTIVE

upvoted 1 times

🗳️ 👤 **Scunningham99** 3 years, 6 months ago

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

upvoted 3 times

🗳️ 👤 **Aesthet** 3 years, 7 months ago

C final answer

upvoted 1 times

🗳️ 👤 **Chhotu_DBA** 3 years, 7 months ago

C is the right one

upvoted 1 times

To meet new data compliance requirements, a company needs to keep critical data durably stored and readily accessible for 7 years. Data that is more than 1 year old is considered archival data and must automatically be moved out of the Amazon Aurora MySQL DB cluster every week. On average, around 10 GB of new data is added to the database every month. A database specialist must choose the most operationally efficient solution to migrate the archival data to Amazon S3.

Which solution meets these requirements?

- A. Create a custom script that exports archival data from the DB cluster to Amazon S3 using a SQL view, then deletes the archival data from the DB cluster. Launch an Amazon EC2 instance with a weekly cron job to execute the custom script.
- B. Configure an AWS Lambda function that exports archival data from the DB cluster to Amazon S3 using a `SELECT INTO OUTFILE S3` statement, then deletes the archival data from the DB cluster. Schedule the Lambda function to run weekly using Amazon EventBridge (Amazon CloudWatch Events).
- C. Configure two AWS Lambda functions: one that exports archival data from the DB cluster to Amazon S3 using the `mysqldump` utility, and another that deletes the archival data from the DB cluster. Schedule both Lambda functions to run weekly using Amazon EventBridge (Amazon CloudWatch Events).
- D. Use AWS Database Migration Service (AWS DMS) to continually export the archival data from the DB cluster to Amazon S3. Configure an AWS Data Pipeline process to run weekly that executes a custom SQL script to delete the archival data from the DB cluster.

Suggested Answer: C

Community vote distribution

B (100%)

🗳️ 👤 **Jaypdv** Highly Voted 3 years, 7 months ago

Going for B. since `SELECT INTO OUTFILE S3` is available on Aurora.

Option C. uses `mysqldump` who does not dump directly to S3

upvoted 14 times

🗳️ 👤 **MultiAZ** Most Recent 1 year, 4 months ago

Selected Answer: B

Answer is B. The data should be readily accessible (e.g. via Athena), so `mysqldump` is not useful

upvoted 1 times

🗳️ 👤 **Pranava_GCP** 1 year, 8 months ago

Selected Answer: B

B. Configure an AWS Lambda function that exports archival data from the DB cluster to Amazon S3 using a `SELECT INTO OUTFILE S3` statement, then deletes the archival data from the DB cluster. Schedule the Lambda function to run weekly using Amazon EventBridge (Amazon CloudWatch Events).

upvoted 1 times

🗳️ 👤 **rag1482** 2 years, 6 months ago

If the amount of data to be selected is large (more than 25 GB), we recommend that you use multiple `SELECT INTO OUTFILE S3` statements to save the data to Amazon S3

Answer: B

upvoted 1 times

🗳️ 👤 **sachin** 2 years, 10 months ago

B is correct approach.

C `mysqldump` can not dump into S3

<https://aws.amazon.com/blogs/database/best-practices-for-exporting-and-importing-data-from-amazon-aurora-mysql-to-amazon-s3/>

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: B

B because:

1. Lambda function max run time is 15 min

<https://aws.amazon.com/about-aws/whats-new/2018/10/aws-lambda-supports-functions-that-can-run-up-to-15-minutes/>

2. SELECT INTO OUTFILES3 is there, and 10GB data per week sounds reasonable to finish copying within 15 min
<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Integrating.SaveIntoS3.html>

AWS DMS can copy to S3 but option D says continually export, while we need weekly

<https://aws.amazon.com/blogs/database/archiving-data-from-relational-databases-to-amazon-glacier-via-aws-dms/>
upvoted 1 times

🗳️ 👤 **pcpcpc888** 3 years, 3 months ago

running a continually DMS job would NOT be operationally efficient, when talk about which, serverless options combined with Lambda and EventBridge would be a much better choice; considering the volume of the weekly archival, the duration would not hit Lambda timeout; however, it seems like more development would be needed for C, cause Select Into Outfile S3 directly integrates with S3. So B.

upvoted 3 times

🗳️ 👤 **Raj12131** 3 years, 4 months ago

Option A requires more effort and hence can be ruled out. Option B uses same lambda function for data migration and deletion thereafter. It doesn't work as lambda might timeout. Option C uses mysqldump which is ok but not as efficient as DMS. Option D is the correct solution in my view.

upvoted 1 times

🗳️ 👤 **Shunpin** 3 years, 5 months ago

Selected Answer: B

For option D, I will consider how DMS export data export to S3 looks like and also how DMS handle "delete" CDC statements. With DMS option, you need additional tasks to filter data and not easy to maintain.

upvoted 1 times

🗳️ 👤 **SMAZ** 3 years, 5 months ago

I think its D

<https://aws.amazon.com/blogs/database/archiving-data-from-relational-databases-to-amazon-glacier-via-aws-dms/>
upvoted 1 times

🗳️ 👤 **jove** 3 years, 5 months ago

Lambda functions have 15mins max execution time. If the extract and delete takes longer than 15 mins using a Lambda function won't work. This limitation might rule out option B and C. Option D will work but "continually export the archival data" is not a requirement.

Thoughts?

upvoted 1 times

🗳️ 👤 **VPup** 3 years, 3 months ago

Good catch on the 15 min limit for Lambda! But in the context of the question - " Each month, around 10 GB of fresh data is uploaded to the database." - I would assume 2.5 GB weekly data volume - seems reasonable to assume that the export and delete will be done within 15 min. so B is still an option here

upvoted 1 times

🗳️ 👤 **Aesthet** 3 years, 7 months ago

B final answer

upvoted 2 times

🗳️ 👤 **manan728** 3 years, 7 months ago

B is correct.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Integrating.SaveIntoS3.html>
upvoted 4 times

🗳️ 👤 **Chhotu_DBA** 3 years, 7 months ago

Option B correct

upvoted 2 times

🗳️ 👤 **novak18** 3 years, 8 months ago

Answer should be D?



https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.S3.html

<https://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/dp-object-sqlactivity.html>



upvoted 1 times

🗳️ 👤 **faramawi** 3 years, 7 months ago

I think it should be D too. I think it provides "most operationally efficient solution to migrate the archival data to Amazon S3 ".
<https://aws.amazon.com/blogs/database/archiving-data-from-relational-databases-to-amazon-glacier-via-aws-dms/>
<https://aws.amazon.com/blogs/database/replicate-data-from-amazon-aurora-to-amazon-s3-with-aws-database-migration-service/>
upvoted 1 times

  **Justu** 3 years, 5 months ago

Can you use AWS Data Pipeline process Custom SQL Query to delete data from RDS?
upvoted 1 times

  **jove** 3 years, 5 months ago

Yes you can but I'm not sure if using DMS is the right option
upvoted 1 times

  **Jiang_aws1** 2 years, 7 months ago

DMS is for DB migration tools & very \$\$\$ so we just use time by time but let it run as job tools. so Lambda is right tools for this .
upvoted 1 times

A company developed a new application that is deployed on Amazon EC2 instances behind an Application Load Balancer. The EC2 instances use the security group named sg-application-servers. The company needs a database to store the data from the application and decides to use an Amazon RDS for MySQL DB instance. The DB instance is deployed in a private DB subnet. What is the MOST restrictive configuration for the DB instance security group?

- A. Only allow incoming traffic from the sg-application-servers security group on port 3306.
- B. Only allow incoming traffic from the sg-application-servers security group on port 443.
- C. Only allow incoming traffic from the subnet of the application servers on port 3306.
- D. Only allow incoming traffic from the subnet of the application servers on port 443.

Suggested Answer: B

Community vote distribution

A (100%)

🗳️ 👤 **shantest1** Highly Voted 3 years, 8 months ago
A. Answer

Database port 3306 and better to allow only the specific subnet instead of the entire subnet.
upvoted 14 times

🗳️ 👤 **novice_expert** 3 years ago
Do we allow traffic from security group
OR
from resources that are assigned to the same security group ?
upvoted 1 times

🗳️ 👤 **Pranava_GCP** Most Recent 1 year, 9 months ago
Selected Answer: A

A. Only allow incoming traffic from the sg-application-servers security group on port 3306.
upvoted 1 times

🗳️ 👤 **Pranava_GCP** 1 year, 9 months ago
Because 1) port 3306 is default port number for mySQL 2) a security group has to be explicitly assigned to an EC2 instance.
upvoted 1 times

🗳️ 👤 **aviathor** 2 years ago
Selected Answer: A
Where on earth do these "Correct answers" come from?

Allowing connections only from members of sg-application-servers is more restrictive than allowing traffic from the whole subnet. 3306 is probably the correct port for RDS. Therefore A
upvoted 2 times

🗳️ 👤 **Bernardes** 2 years, 6 months ago
I'll go with C:
Security groups contains rules allowing or denying access to specified IP address and TCP Ports. Then they are associated with resources (such as ec2, rds, etc).
The question says that EC2 use the security group sg-application-servers. That is the sg-application-servers contains the rules that at this moment make the ec2 communication work. To allow the RDS instance talk with this ec2, is necessary to create a SG and specify the address of the ec2 instances in the ingress rules, referencing the port 3306. Or, allow the traffic from the entire subnet at this same port.
I've read the security group doc again and dont see nothing about grouping aws resources to reference as some kind of "security resource group" as the A answer say.

https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html
upvoted 1 times

🗨️ 👤 **im_not_robot** 2 years, 3 months ago

C is wrong. Because allow the whole subnets mean allow all application on that subnet --> not secure.

FYI: security group doesn't have 'Deny' rule.

A is the answer. Security A allow traffic from Security B mean that Security A allow all resources using Security B
upvoted 1 times

🗨️ 👤 **novice_expert** 3 years ago

Selected Answer: A

A. Only allow incoming traffic from the sg-application-servers security group on port 3306.

upvoted 3 times

🗨️ 👤 **awsmonster** 3 years, 4 months ago

Ans: A

upvoted 1 times

🗨️ 👤 **Hits_23** 3 years, 6 months ago

Answer A. most restrictive approach is to allow only incoming connections from SG of EC2 instance on port 3306

upvoted 1 times

🗨️ 👤 **Suresh108** 3 years, 7 months ago

AAAAAAAAAAAAA

upvoted 2 times

🗨️ 👤 **Aesthet** 3 years, 7 months ago

A final answer

upvoted 1 times

A company is moving its fraud detection application from on premises to the AWS Cloud and is using Amazon Neptune for data storage. The company has set up a 1 Gbps AWS Direct Connect connection to migrate 25 TB of fraud detection data from the on-premises data center to a Neptune DB instance. The company already has an Amazon S3 bucket and an S3 VPC endpoint, and 80% of the company's network bandwidth is available.

How should the company perform this data load?

- A. Use an AWS SDK with a multipart upload to transfer the data from on premises to the S3 bucket. Use the Copy command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.
- B. Use AWS Database Migration Service (AWS DMS) to transfer the data from on premises to the S3 bucket. Use the Loader command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.
- C. Use AWS DataSync to transfer the data from on premises to the S3 bucket. Use the Loader command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.
- D. Use the AWS CLI to transfer the data from on premises to the S3 bucket. Use the Copy command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.

Suggested Answer: C


Community vote distribution

C (100%)

 **Jaypdy** Highly Voted 3 years, 7 months ago

Answer is C. since DMS can only use databases (and S3) as sources, and the question does not specify that the on-prem data resides in a DB. In which case, Datasync is a more likely choice.

upvoted 8 times

 **manan728** Highly Voted 3 years, 6 months ago

This question was asked in my exam. I went with C.

upvoted 6 times

 **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: C

C. Use AWS DataSync to transfer the data from on premises to the S3 bucket. Use the Loader command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.

AWS DataSync is used to transfer data from On-prem to S3, no source db is mentioned.

To load data from S3 to AWS Neptune to use loader command

<https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html>

<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load.html>

upvoted 1 times

 **novice_expert** 3 years ago

Selected Answer: C

DMS target can not be S3

AWS DataSync is a secure, online service that automates and accelerates moving data between on-premises and AWS storage services.

C. Use AWS DataSync to transfer the data from on premises to the S3 bucket. Use the Loader command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.

aws DataSync is the way to transfer data from on-prem to S3, no source db is mentioned.

Only way to load data from S3 to AWS Neptune is using loader

<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load.html>

upvoted 2 times

🗨️ 👤 **awsjji** 2 years, 7 months ago

Answer is C. But DMS target can be S3. https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.S3.html

upvoted 1 times

🗨️ 👤 **RotterDam** 3 years, 2 months ago

Got this question in my exam. (i cleared it). C is correct

upvoted 2 times

🗨️ 👤 **awsmonster** 3 years, 4 months ago

Answer should be B, use DMS.

Data Sync does not read from databases (<https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html>).

"AWS DataSync is an online data transfer service that simplifies, automates, and accelerates moving data between on-premises storage systems and AWS storage services, and also between AWS storage services."

upvoted 1 times

🗨️ 👤 **awsmonster** 3 years, 4 months ago

Sorry .. Answer is C.

My error, the questions says move fraud data, nothing is mentioned about database.

upvoted 1 times

🗨️ 👤 **Hits_23** 3 years, 6 months ago

C is correct answer.

aws DataSync is the way to transfer data from on-prem to S3, no source db is mentioned.

Only way to load data from S3 to AWS Neptune is using loader

<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load.html>

upvoted 3 times

🗨️ 👤 **Aesthet** 3 years, 6 months ago

C is the answer

If on-premise DB was mentioned I would choose B

upvoted 1 times

🗨️ 👤 **db_interest** 3 years, 6 months ago

C makes sense

upvoted 1 times

🗨️ 👤 **manan728** 3 years, 7 months ago

Between B and C that are the only plausible options C seems more legit for the given scenario.

upvoted 1 times

🗨️ 👤 **novak18** 3 years, 7 months ago

Answer should be B I guess

upvoted 1 times

A company migrated one of its business-critical database workloads to an Amazon Aurora Multi-AZ DB cluster. The company requires a very low RTO and needs to improve the application recovery time after database failovers.

Which approach meets these requirements?

- A. Set the max_connections parameter to 16,000 in the instance-level parameter group.
- B. Modify the client connection timeout to 300 seconds.
- C. Create an Amazon RDS Proxy database proxy and update client connections to point to the proxy endpoint.
- D. Enable the query cache at the instance level.

Suggested Answer: D

Community vote distribution


C (100%)

 **manan728**  3 years, 6 months ago

C looks to be the winner.

Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. With RDS Proxy, failover times for Aurora and RDS databases are reduced by up to 66% and database credentials, authentication, and access can be managed through integration with AWS Secrets Manager and AWS Identity and Access Management (IAM).


upvoted 15 times

 **novak18**  3 years, 8 months ago

Can the answer be C?

<https://aws.amazon.com/rds/proxy/>

upvoted 6 times

 **Jaypdy** 3 years, 7 months ago

I'm sorry, you're correct. Answer is C.

upvoted 1 times

 **Pranava_GCP**  1 year, 8 months ago

C. Amazon RDS Proxy database proxy

"Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. With RDS Proxy, failover times for Aurora and RDS databases are reduced by up to 66% and database credentials, authentication, and access can be managed through integration with AWS Secrets Manager and AWS Identity and Access Management (IAM)."

<https://aws.amazon.com/rds/proxy/>

upvoted 1 times

 **KaranGandhi30** 3 years ago

Selected Answer: C

When your primary instance is down ABD won't work.

upvoted 1 times

 **novice_expert** 3 years ago

Selected Answer: C

<https://aws.amazon.com/rds/proxy/>

upvoted 1 times

 **RotterDam** 3 years, 2 months ago

Selected Answer: C

Definitely C. Proxies maintain a pool of client connections that respond to failovers and actually IMPROVE failover times by 66%

<https://aws.amazon.com/rds/proxy/>

upvoted 2 times

 **user0001** 3 years, 2 months ago



i vote for C,
the key is "application recovery time after database failovers"
upvoted 1 times

  **ChauPhan** 3 years, 6 months ago



C is correct. A D is not relevant, B even makes failover time slower.
upvoted 1 times

  **Aesthet** 3 years, 6 months ago

C final answer
upvoted 2 times

  **Jaypdy** 3 years, 7 months ago

B. Answer
upvoted 1 times

  **Jaypdy** 3 years, 7 months ago

Oops. Answer is C.
upvoted 1 times

A company is using an Amazon RDS for MySQL DB instance for its internal applications. A security audit shows that the DB instance is not encrypted at rest. The company's application team needs to encrypt the DB instance. What should the team do to meet this requirement?

- A. Stop the DB instance and modify it to enable encryption. Apply this setting immediately without waiting for the next scheduled RDS maintenance window.
- B. Stop the DB instance and create an encrypted snapshot. Restore the encrypted snapshot to a new encrypted DB instance. Delete the original DB instance, and update the applications to point to the new encrypted DB instance.
- C. Stop the DB instance and create a snapshot. Copy the snapshot into another encrypted snapshot. Restore the encrypted snapshot to a new encrypted DB instance. Delete the original DB instance, and update the applications to point to the new encrypted DB instance.
- D. Create an encrypted read replica of the DB instance. Promote the read replica to master. Delete the original DB instance, and update the applications to point to the new encrypted DB instance.

Suggested Answer: A

Reference:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html>

Community vote distribution

C (100%)

🗳️ 👤 **Jaypdr** Highly Voted 3 years, 7 months ago

C. Answer
upvoted 8 times

🗳️ 👤 **shantest1** Highly Voted 3 years, 7 months ago

C. Answer
upvoted 5 times

🗳️ 👤 **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: C
C is correct. You can only enable encryption during copy of the snapshot or when you create DB instance .
upvoted 2 times

🗳️ 👤 **sirfans** 2 years, 7 months ago

Selected Answer: C
C is the right choice. <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html>
upvoted 2 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: C
C. Stop the DB instance and create a snapshot. Copy the snapshot into another encrypted snapshot. Restore the encrypted snapshot to a new encrypted DB instance. Delete the original DB instance, and update the applications to point to the new encrypted DB instance.
upvoted 3 times

🗳️ 👤 **kret** 3 years, 1 month ago

Selected Answer: C
C is the way. You can only enable encryption during copy operation of the snapshot.
upvoted 4 times

🗳️ 👤 **ChauPhan** 3 years, 6 months ago

C is correct
upvoted 1 times

🗳️ 👤 **Aesthet** 3 years, 7 months ago

C final answer
upvoted 1 times

🗳️ 👤 **manan728** 3 years, 7 months ago

Deleting the old database was kind of an unnecessary step added there.

upvoted 1 times

  **manan728** 3 years, 7 months ago

Eh, maybe required for compliance I suppose but C it is.

upvoted 1 times

  **novak18** 3 years, 7 months ago

Answer should be C

upvoted 2 times

A database specialist must create nightly backups of an Amazon DynamoDB table in a mission-critical workload as part of a disaster recovery strategy.

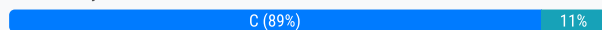
Which backup methodology should the database specialist use to MINIMIZE management overhead?

- A. Install the AWS CLI on an Amazon EC2 instance. Write a CLI command that creates a backup of the DynamoDB table. Create a scheduled job or task that runs the command on a nightly basis.
- B. Create an AWS Lambda function that creates a backup of the DynamoDB table. Create an Amazon CloudWatch Events rule that runs the Lambda function on a nightly basis.
- C. Create a backup plan using AWS Backup, specify a backup frequency of every 24 hours, and give the plan a nightly backup window.
- D. Configure DynamoDB backup and restore for an on-demand backup frequency of every 24 hours.

Suggested Answer: D

On-demand backup allows you to create full backups of your Amazon DynamoDB table for data archiving, helping you meet your corporate and governmental regulatory requirements. You can back up tables from a few megabytes to hundreds of terabytes of data, with no impact on performance and availability to your production applications. Backups process in seconds regardless of the size of your tables, so you do not have to worry about backup schedules or long-running processes. In addition, all backups are automatically encrypted, cataloged, easily discoverable, and retained until explicitly deleted.

Community vote distribution



🗳️ 👤 **KikiNoviandi** 1 year, 7 months ago

Selected Answer: C

c is correct

upvoted 2 times

🗳️ 👤 **shammous** 2 years, 9 months ago

Selected Answer: C

The key word here is "minimum overhead". Using AWS Backup would insure that.

Ref: <https://aws.amazon.com/blogs/database/set-up-scheduled-backups-for-amazon-dynamodb-using-aws-backup/>

upvoted 3 times

🗳️ 👤 **Omijh** 2 years, 11 months ago

Selected Answer: C

C is correct, If you go in the Dynamodb console and click scheduled backup it just takes you to aws backup to create a plan. Since it's nightly scheduled backups that's needed then aws back. If they wanted ondemand manual backup then Ondemand in Dynamodb

upvoted 2 times

🗳️ 👤 **sachin** 2 years, 11 months ago

C is correct as it is using backup service. D option can only take manual backup.

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: D

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

upvoted 2 times

🗳️ 👤 **novice_expert** 3 years ago

correction: Ans is C

use AWS backup service

upvoted 1 times

🗳️ 👤 **Dantas** 3 years, 2 months ago

Selected Answer: C

C is better than D, as the latter introduces administration overhead: Amazon DynamoDB supports stand-alone on-demand backup and restores features. You can create table backups using the console, the AWS Command Line Interface (AWS CLI), or the DynamoDB API.

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

upvoted 3 times

🗨️ 👤 **user0001** 3 years, 2 months ago

Answer is C.

i'm not sure why they are providing the wrong answers for most of the questions

upvoted 1 times

🗨️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: C

agree with other comments

upvoted 1 times

🗨️ 👤 **user0001** 3 years, 3 months ago

C

If you don't want to create scheduling scripts and cleanup jobs, you can use AWS Backup to create backup plans with schedules and retention policies for your DynamoDB tables. AWS Backup runs the backups and deletes them when they expire. For more information, see the AWS Backup Developer Guide.

upvoted 1 times

🗨️ 👤 **Hariru** 3 years, 3 months ago

Why not D

upvoted 1 times

🗨️ 👤 **soyyodario** 3 years, 3 months ago

Selected Answer: C

Answer is C

To be compliance with disaster recovery it's necessary to use AWS Backup: (link bellow) To create backup copies across AWS accounts and Regions and for other advanced features, you should use AWS Backup.

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

upvoted 3 times

🗨️ 👤 **peacegrace** 3 years, 4 months ago

Selected Answer: C

C seems more appropriate, as it has more control

upvoted 1 times

🗨️ 👤 **mnzsql365** 3 years, 5 months ago

Ans is C more control

upvoted 1 times

🗨️ 👤 **akiraklaus** 3 years, 5 months ago

C and D do the same thing, more D reduce overhead.

answer is D

upvoted 1 times

🗨️ 👤 **shuraosipov** 3 years, 5 months ago

Selected Answer: C

Answer is C.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/CreateBackup.html#:~:text=If%20you%20don%27t%20want%20to%20create%20sch>

upvoted 2 times

🗨️ 👤 **damaldon** 3 years, 7 months ago

Answer is C, as per the following link:

If you don't want to create scheduling scripts and cleanup jobs, you can use AWS Backup to create backup plans with schedules and retention policies for your DynamoDB tables. AWS Backup runs the backups and deletes them when they expire

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

upvoted 3 times

🗨️ 👤 **Aesthet** 3 years, 7 months ago

C final answer

upvoted 3 times

A company is using a Single-AZ Amazon RDS for MySQL DB instance for development. The DB instance is experiencing slow performance when queries run.

Amazon CloudWatch metrics indicate that the instance requires more I/O capacity.

Which actions can a database specialist perform to resolve this issue? (Choose two.)

- A. Restart the application tool used to run queries.
- B. Change to a database instance class with higher throughput.
- C. Convert from Single-AZ to Multi-AZ.
- D. Increase the I/O parameter in Amazon RDS Enhanced Monitoring.
- E. Convert from General Purpose to Provisioned IOPS (PIOPS).

Suggested Answer: BD

Community vote distribution

BE (100%)

🗳️ 👤 **lollyj** 2 years, 5 months ago

Selected Answer: BE

My choices...

upvoted 3 times

🗳️ 👤 **rags1482** 2 years, 6 months ago

The I/O capacity of the instance is based on the instance storage type and size.

So Answer is B & E

upvoted 1 times

🗳️ 👤 **Adi_M** 2 years, 8 months ago

Selected Answer: BE

D. Increase the I/O parameter in Amazon RDS Enhanced Monitoring. - not the right answer. Just enabling enhanced monitoring is not going to help the performance.

correct Answer: B & E

upvoted 2 times

🗳️ 👤 **xyj** 2 years, 8 months ago

B & E

<https://aws.amazon.com/blogs/database/best-storage-practices-for-running-production-workloads-on-hosted-databases-with-amazon-rds-or-amazon-ec2/>

upvoted 3 times

A company has an AWS CloudFormation template written in JSON that is used to launch new Amazon RDS for MySQL DB instances. The security team has asked a database specialist to ensure that the master password is automatically rotated every 30 days for all new DB instances that are launched using the template.

What is the MOST operationally efficient solution to meet these requirements?

- A. Save the password in an Amazon S3 object. Encrypt the S3 object with an AWS KMS key. Set the KMS key to be rotated every 30 days by setting the EnableKeyRotation property to true. Use a CloudFormation custom resource to read the S3 object to extract the password.
- B. Create an AWS Lambda function to rotate the secret. Modify the CloudFormation template to add an AWS::SecretsManager::RotationSchedule resource. Configure the RotationLambdaARN value and, for the RotationRules property, set the AutomaticallyAfterDays parameter to 30.
- C. Modify the CloudFormation template to use the AWS KMS key as the database password. Configure an Amazon EventBridge rule to invoke the KMS API to rotate the key every 30 days by setting the ScheduleExpression parameter to `*/30/*`.
- D. Integrate the Amazon RDS for MySQL DB instances with AWS IAM and centrally manage the master database user password.

Suggested Answer: C

Community vote distribution

B (100%)

 **Jaypdv** Highly Voted 3 years, 7 months ago


B. Answer

upvoted 12 times

 **khun** Most Recent 2 years, 5 months ago

B is the answer. AWS secret supports rotation

upvoted 2 times

 **Bobapo2** 2 years, 12 months ago

Selected Answer: B

B. Answer

upvoted 1 times

 **novice_expert** 3 years ago

Selected Answer: B

Secrete Manager -> Lambda to rotate secret -> modify Cloud formation to add rotation schedule

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-secretsmanager-rotationschedule.html>

upvoted 2 times

 **Dantas** 3 years, 2 months ago

Selected Answer: B

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-secretsmanager-rotationschedule.html>


upvoted 1 times

 **RotterDam** 3 years, 2 months ago

Selected Answer: B

(B) is the correct answer. (Who is making the official chosen answers? Almost all of them are wrong - the community ones are the correct ones)


upvoted 4 times

 **tugboat** 3 years, 3 months ago

Selected Answer: B

Lambda with Secrets Manager works perfectly

upvoted 1 times

 **kped21** 3 years, 3 months ago

B - Lambda with secretsManager.

upvoted 1 times

 **Aesthet** 3 years, 6 months ago

B final answer

upvoted 2 times

🗨️ 👤 **manan728** 3 years, 7 months ago

This question was asked in my exam. B is correct.

upvoted 2 times

🗨️ 👤 **novak18** 3 years, 8 months ago

Shouldn't the answer be B?

upvoted 1 times

🗨️ 👤 **novak18** 3 years, 7 months ago

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-secretsmanager-rotationschedule.html>

upvoted 3 times

A startup company is building a new application to allow users to visualize their on-premises and cloud networking components. The company expects billions of components to be stored and requires responses in milliseconds. The application should be able to identify:

- ⇒ The networks and routes affected if a particular component fails.
- ⇒ The networks that have redundant routes between them.
- ⇒ The networks that do not have redundant routes between them.
- ⇒ The fastest path between two networks.

Which database engine meets these requirements?

- A. Amazon Aurora MySQL
- B. Amazon Neptune
- C. Amazon ElastiCache for Redis
- D. Amazon DynamoDB

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **novak18** Highly Voted 3 years, 7 months ago
Answer is B
upvoted 6 times

🗳️ 👤 **Kronos369** Most Recent 1 year, 1 month ago
Selected Answer: B
The use cases show that its a graph db
upvoted 1 times

🗳️ 👤 **Kodoma** 1 year, 11 months ago
Selected Answer: B
Graphs = Neptune (B!)
upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago
Selected Answer: B
B. Amazon Neptune
upvoted 1 times

🗳️ 👤 **Suresh108** 3 years, 7 months ago
BBBBBBBBBB

relationship based store is Neptune.
upvoted 1 times

🗳️ 👤 **Aesthet** 3 years, 7 months ago
B final answer
upvoted 4 times

🗳️ 👤 **manan728** 3 years, 7 months ago
This question was asked in my exam. Looks like aws do not have many variety of questions on Neptune.
upvoted 3 times

🗳️ 👤 **shantest1** 3 years, 7 months ago
B. answer
upvoted 2 times

An online retail company is planning a multi-day flash sale that must support processing of up to 5,000 orders per second. The number of orders and exact schedule for the sale will vary each day. During the sale, approximately 10,000 concurrent users will look at the deals before buying items. Outside of the sale, the traffic volume is very low. The acceptable performance for read/write queries should be under 25 ms. Order items are about 2 KB in size and have a unique identifier. The company requires the most cost-effective solution that will automatically scale and is highly available.

Which solution meets these requirements?

- A. Amazon DynamoDB with on-demand capacity mode
- B. Amazon Aurora with one writer node and an Aurora Replica with the parallel query feature enabled
- C. Amazon DynamoDB with provisioned capacity mode with 5,000 write capacity units (WCUs) and 10,000 read capacity units (RCUs)
- D. Amazon Aurora with one writer node and two cross-Region Aurora Replicas

Suggested Answer: C

Community vote distribution



A (100%)

  **Jaypdv** Highly Voted 3 years, 8 months ago

I'll go for A.

I think C. is a trap; if you are writing 5000 items at 2KB/second then you need 10000 WCU and 5000 RCU (assuming eventually consistent reads). C. has it in reverse.

upvoted 14 times

  **Jaypdv** 3 years, 7 months ago

And C. does not automatically scale as in the requirement

upvoted 3 times

  **shantest1** 3 years, 7 months ago

A seems to be correct

I think so as well, c is in reverse

upvoted 1 times

  **manan728** 3 years, 7 months ago

For 5000 RCU you are assuming strongly consistent reads. If it's an eventual consistent read you only need 2500 RCU



upvoted 2 times

  **Sathish_dbs** Most Recent 1 year, 7 months ago

Selected Answer: A

the next closest answer would be DynamoDB with Auto scaling but that it not part of the answers hence on-demand is the best answer

upvoted 1 times

  **IhorK** 1 year, 9 months ago

Selected Answer: A

With 2KB size for 5,000 orders per second we need 10 000 WCU (1 table write/sec = 1 WCU with blocks of 1KB). C suggest 5,000 write capacity units (WCUs), not enough.

upvoted 1 times

  **awsjji** 2 years, 7 months ago

Selected Answer: A

1250 RCUs can read 10,000 KB. A is correct

upvoted 1 times

  **awsjji** 2 years, 7 months ago

1250 RCUs can read 10,000 KB if its eventual consistent read. A is correct

upvoted 1 times

  **awsjji** 2 years, 7 months ago

2500 RCU for reads

upvoted 2 times

🗨️ 👤 **praffuln** 3 years ago

Selected Answer: A

I'll go for A. C is not scalable & cost affective.

upvoted 1 times

🗨️ 👤 **novice_expert** 3 years ago

Selected Answer: A

25 milliseconds

Outside of the sale, the amount of traffic is really minimal

The business seeks the most cost-effective solution possible that is both highly accessible and scales automatically.

upvoted 1 times

🗨️ 👤 **jove** 3 years, 5 months ago

Selected Answer: A

It is A

upvoted 3 times

🗨️ 👤 **GMartinelli** 3 years, 5 months ago

Selected Answer: A

Option A

upvoted 4 times

🗨️ 👤 **ChauPhan** 3 years, 6 months ago

he number of orders and exact schedule for the sale will vary each day. During the sale, approximately 10,000 concurrent users will look at the deals before buying items. Outside of the sale, the traffic volume is very low

=> Setting provisioning DynamoDB fix read 5000/write 10000 with will waste the resource when the traffic is low. It is not cost-effective.

I go with A

upvoted 3 times

🗨️ 👤 **Aesthet** 3 years, 6 months ago

A final answer

upvoted 2 times

🗨️ 👤 **Zhongkai** 3 years, 7 months ago

c does not mention auto scaling. in addition, 10000 WCU are needed. I will go with A

upvoted 3 times

🗨️ 👤 **db_interest** 3 years, 7 months ago

C seems correct

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

upvoted 1 times

🗨️ 👤 **db_interest** 3 years, 6 months ago

Correction. A is the answer

upvoted 1 times

🗨️ 👤 **novak18** 3 years, 8 months ago

Answer is C

upvoted 1 times

A ride-hailing application uses an Amazon RDS for MySQL DB instance as persistent storage for bookings. This application is very popular and the company expects a tenfold increase in the user base in next few months. The application experiences more traffic during the morning and evening hours.

This application has two parts:

- ⇒ An in-house booking component that accepts online bookings that directly correspond to simultaneous requests from users.
- ⇒ A third-party customer relationship management (CRM) component used by customer care representatives. The CRM uses queries to access booking data.

A database specialist needs to design a cost-effective database solution to handle this workload.

Which solution meets these requirements?

- A. Use Amazon ElastiCache for Redis to accept the bookings. Associate an AWS Lambda function to capture changes and push the booking data to the RDS for MySQL DB instance used by the CRM.
- B. Use Amazon DynamoDB to accept the bookings. Enable DynamoDB Streams and associate an AWS Lambda function to capture changes and push the booking data to an Amazon SQS queue. This triggers another Lambda function that pulls data from Amazon SQS and writes it to the RDS for MySQL DB instance used by the CRM.
- C. Use Amazon ElastiCache for Redis to accept the bookings. Associate an AWS Lambda function to capture changes and push the booking data to an Amazon Redshift database used by the CRM.
- D. Use Amazon DynamoDB to accept the bookings. Enable DynamoDB Streams and associate an AWS Lambda function to capture changes and push the booking data to Amazon Athena, which is used by the CRM.

Suggested Answer: A

Community vote distribution

B (100%)

🗳️ 👤 **Aesthet** Highly Voted 3 years, 7 months ago

In A "AWS Lambda function to capture changes" capture changes to what? ElastiCache? The main use of ElastiCache is to cache frequently read data. Also "the company expects a tenfold increase in the user base" and "correspond to simultaneous requests from users" suggest DynamoDB imo. Since we already have MySQL RDS I prefer answer B
upvoted 11 times

🗳️ 👤 **learnazureportal** 3 years, 7 months ago

Actually, It is D. based on the increasing the overall, we cannot go for B.
upvoted 4 times

🗳️ 👤 **ChauPhan** 3 years, 7 months ago

Amazon Athena is an interactive query service that makes it easy to analyze data directly in Amazon Simple Storage Service (Amazon S3) using standard SQL.

It is DB query service serverless for data stored on S3. It is itself not a DB storage
upvoted 4 times

🗳️ 👤 **cnmc** 2 years ago

Athena can read from DynamoDB... This kind of answers show that you guys are only studying for the cert without any actual experience. Setting up a whole new RDS clusters just so some CRM can read from it will cost your company thousands of dollars per month.
upvoted 1 times

🗳️ 👤 **JasonZhu** 1 year, 7 months ago

Athena is more cost-effective than RDS.
I prefer D.
upvoted 1 times



🗳️ 👤 **manan728** Highly Voted 3 years, 7 months ago

This question wasn't asked but in summary a lot of questions exactly from this 145 set questions were asked. I passed with a score of 794. Bunch of new questions on Aurora upgrades and maintenance windows, MS SQL server and Oracle migrations. All the best 🍀
upvoted 7 times

🗳️ 👤 **user0001** 3 years, 2 months ago



you are trying to mention score to give gradability to your answer, you only get pass or fail but not the score

upvoted 1 times

  **guru_ji** 3 years, 6 months ago

I gave exam today. Only 60% questions came in actual exam from this 145 set.

upvoted 2 times

  **palanikumar_n** 3 years, 6 months ago

Do you have the 145 questions. I could see only 112 question in the site

upvoted 1 times


  **aviathor** Most Recent 1 year, 11 months ago

Selected Answer: B

Although I do not understand why one should use one Lambda function to read from the DynamoDB stream, write to an SQS queue and then another Lambda to read from SQS and write to RDS... B is what makes most sense.

But I would have considered using only one Lambda function and no SQS.



upvoted 1 times

  **manig** 2 years, 1 month ago

I go with Answer D.

For booking we need a persistent DB which is a dynamoDB but querying the data through Athena is more cost effective way.

upvoted 1 times

  **lollyj** 2 years, 5 months ago

Selected Answer: B

Chose this option because this architecture definitely needs a SQS component being that it's a ride hailing application. Only makes sense

upvoted 1 times



  **novice_expert** 3 years ago

Selected Answer: B

no use of ElastiCache + Redis (A, C)

Athena is not DB (D)

upvoted 1 times

  **Dantas** 3 years, 2 months ago

Selected Answer: B

- Redis is an in-memory data store, therefore not suitable for bookings persistence. A and C are incorrect.



- Athena is an interactive query service, not a database. D is incorrect.

upvoted 4 times

  **user0001** 3 years, 2 months ago

B is the answer , we need the data in mysql

upvoted 1 times

  **tugboat** 3 years, 3 months ago

Selected Answer: B

CRM needs data in MySQL

upvoted 2 times


  **user0001** 3 years, 3 months ago

i go with B

you need to save data to MySQL

A is not persistence

upvoted 2 times

  **scottkerker** 3 years, 5 months ago



ElastiCache, either MemcaheD or Redis, is used for data frequently read. And Amazon Athena is for querying data stored in Amazon S3. Considering all the features provided by the services, Option B is the best option.

upvoted 1 times

  **jove** 3 years, 6 months ago

Answer is B

upvoted 1 times

  **GMartinelli** 3 years, 6 months ago

Selected Answer: B

It's B

upvoted 4 times

🗨️ 👤 **Amy2009** 3 years, 6 months ago

should be A

upvoted 3 times

🗨️ 👤 **VNKSEC** 3 years, 6 months ago

@guru_ji ? only 60% ? , hm... you can send me email

upvoted 1 times

🗨️ 👤 **guru_ji** 3 years, 6 months ago

I gave exam today. Only 60% questions came in actual exam.

upvoted 1 times

🗨️ 👤 **guru_ji** 3 years, 7 months ago

Answer: B

upvoted 1 times

An online advertising website uses an Amazon DynamoDB table with on-demand capacity mode as its data store. The website also has a DynamoDB Accelerator (DAX) cluster in the same VPC as its web application server. The application needs to perform infrequent writes and many strongly consistent reads from the data store by querying the DAX cluster.

During a performance audit, a systems administrator notices that the application can look up items by using the DAX cluster. However, the QueryCacheHits metric for the DAX cluster consistently shows 0 while the QueryCacheMisses metric continuously keeps growing in Amazon CloudWatch.

What is the MOST likely reason for this occurrence?

- A. A VPC endpoint was not added to access DynamoDB.
- B. Strongly consistent reads are always passed through DAX to DynamoDB.
- C. DynamoDB is scaling due to a burst in traffic, resulting in degraded performance.
- D. A VPC endpoint was not added to access CloudWatch.


Suggested Answer: B

Reference:

<https://github.com/aws/aws-sdk-java/issues/1983>

Community vote distribution

B (100%)

 **grekh001** Highly Voted 3 years, 8 months ago

B is correct.

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

"If the request specifies strongly consistent reads, DAX passes the request through to DynamoDB. The results from DynamoDB are not cached in DAX. Instead, they are simply returned to the application."

upvoted 6 times

 **Pranava_GCP** Most Recent 1 year, 8 months ago

Selected Answer: B

B. Strongly consistent reads are always passed through DAX to DynamoDB.

upvoted 1 times

 **lollyj** 2 years, 5 months ago

Selected Answer: B

My thoughts - if strongly consistency reads are requested it will pass through DAX to the DB.

upvoted 1 times

 **novice_expert** 3 years ago

Selected Answer: B

strong consistency => bypass DAX

upvoted 3 times


 **RotterDam** 3 years, 2 months ago

Selected Answer: B

Strongly Consistent Reads are considered PASS THROUGH and will never update DAX

Interestingly - TransactionGetItems is also a PASS THROUGH but TransactionWriteItems is not!!

upvoted 2 times

 **kped21** 3 years, 3 months ago

B:

If the request specifies eventually consistent reads (the default behavior), it tries to read the item from DAX:

If DAX has the item available (a cache hit), DAX returns the item to the application without accessing DynamoDB.

If DAX does not have the item available (a cache miss), DAX passes the request through to DynamoDB. When it receives the response from

DynamoDB, DAX returns the results to the application. But it also writes the results to the cache on the primary node.

If the request specifies strongly consistent reads, DAX passes the request through to DynamoDB. The results from DynamoDB are not cached in DAX. Instead, they are simply returned to the application.

upvoted 3 times

A financial company recently launched a portfolio management solution. The backend of the application is powered by Amazon Aurora with MySQL compatibility.

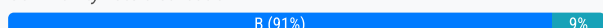
The company requires an RTO of 5 minutes and an RPO of 5 minutes. A database specialist must configure an efficient disaster recovery solution with minimal replication lag.

Which approach should the database specialist take to meet these requirements?

- A. Configure AWS Database Migration Service (AWS DMS) and create a replica in a different AWS Region.
- B. Configure an Amazon Aurora global database and add a different AWS Region.
- C. Configure a binlog and create a replica in a different AWS Region.
- D. Configure a cross-Region read replica.

Suggested Answer: D

Community vote distribution



🗳️ 👤 **[Removed]** 1 year, 11 months ago

Option D could be also correct, but according to documentation:

"The promotion process takes a few minutes to complete. When you promote a read replica, replication is stopped and the DB instances are rebooted. When the reboot is complete, the read replica is available as a new DB cluster."

Therefore it does not suit the condition, that RTO should be less than 5 minutes

upvoted 1 times

🗳️ 👤 **manig** 2 years, 1 month ago

B

Aurora Global cluster RPO typically < 1Mins and RTO depends on how soon we detach the other region cluster from Global cluster and promote to primary. With simple Lambda/jenkins job this can be easily achievable within < 5 Mins.

upvoted 2 times

🗳️ 👤 **OCHT** 2 years, 3 months ago

Selected Answer: B

Verified.

upvoted 1 times

🗳️ 👤 **sachin** 2 years, 11 months ago

B Aurora Global Tables. (As an alternative to cross-Region read replicas, you can scale read operations with minimal lag time by using an Aurora global database.)

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: B

5 min RTO 5 min RPO

Aurora GLOBAL RPO 1second and RTO 1 minute

upvoted 1 times

🗳️ 👤 **Dantas** 3 years, 2 months ago

Selected Answer: B

Amazon Aurora global databases span multiple AWS Regions, enabling low latency global reads and providing fast recovery from the rare outage that might affect an entire AWS Region.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-global-database.html>

<https://www.amazonaws.cn/en/rds/aurora/faqs/>

upvoted 4 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: B

question is vague. what does "The firm demands a response time of five minutes and a response time of five minutes. " mean? I assume they mean RTO and RPO. If 5 minutes is the same for both then Aurora GLOBAL will suffice (RPO 1second and RTO 1 minute)

upvoted 4 times

🗨️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: C

Efficient is key - takes our global DBs

Read-replicas are efficient and low-latency

upvoted 1 times

🗨️ 👤 **BreeRawkus** 3 years, 1 month ago

Global DB uses storage replication = low replication latency, Answer is B

upvoted 2 times

🗨️ 👤 **user0001** 3 years, 3 months ago

Answer B

upvoted 1 times

🗨️ 👤 **mnzsql365** 3 years, 5 months ago

B for me

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-global-database-disaster-recovery.html>

upvoted 4 times

🗨️ 👤 **nood** 3 years, 5 months ago

B for me

<https://aws.amazon.com/blogs/database/how-to-choose-the-best-disaster-recovery-option-for-your-amazon-aurora-mysql-cluster/>

upvoted 2 times

🗨️ 👤 **jove** 3 years, 6 months ago

If cost is not a concern, global database is a better option than read replica.

upvoted 1 times

🗨️ 👤 **grekh001** 3 years, 6 months ago

I think the answer is B

<https://aws.amazon.com/about-aws/whats-new/2019/11/aurora-supports-in-place-conversion-to-global-database/>

<https://aws.amazon.com/blogs/database/how-to-choose-the-best-disaster-recovery-option-for-your-amazon-aurora-mysql-cluster/>

"Aurora Global Database provides the lowest consistent RTO and RPO option while requiring the least management overhead."

upvoted 4 times

🗨️ 👤 **johnconnor** 3 years, 6 months ago

I think you are right, the other options would work but B seems to be the faster

upvoted 3 times

A company hosts an internal file-sharing application running on Amazon EC2 instances in VPC_A. This application is backed by an Amazon ElastiCache cluster, which is in VPC_B and peered with VPC_A. The company migrates its application instances from VPC_A to VPC_B. Logs indicate that the file-sharing application no longer can connect to the ElastiCache cluster.

What should a database specialist do to resolve this issue?

- A. Create a second security group on the EC2 instances. Add an outbound rule to allow traffic from the ElastiCache cluster security group.
- B. Delete the ElastiCache security group. Add an interface VPC endpoint to enable the EC2 instances to connect to the ElastiCache cluster.
- C. Modify the ElastiCache security group by adding outbound rules that allow traffic to VPC_B's CIDR blocks from the ElastiCache cluster.
- D. Modify the ElastiCache security group by adding an inbound rule that allows traffic from the EC2 instances' security group to the ElastiCache cluster.

Suggested Answer: C

Community vote distribution

D (100%)

🗳️ 👤 **SamDDD** 1 year, 10 months ago

Why not B? <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/elasticache-privatelink.html>

upvoted 1 times

🗳️ 👤 **novice_expert** 3 years ago

Selected Answer: D

You can update security groups to reference peer VPC group

<https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-security-groups.html>

upvoted 3 times

🗳️ 👤 **RotterDam** 3 years, 2 months ago

Selected Answer: D

Has to be D

upvoted 3 times

🗳️ 👤 **tugboat** 3 years, 3 months ago

Selected Answer: D

peering is best option

upvoted 2 times

🗳️ 👤 **Hariru** 3 years, 3 months ago

Why not B?

upvoted 1 times

🗳️ 👤 **2025flakyt** 3 years, 5 months ago

D is correct

upvoted 1 times

🗳️ 👤 **jove** 3 years, 6 months ago

Answer is D

upvoted 1 times

🗳️ 👤 **leunamE** 3 years, 6 months ago

Option D.

upvoted 4 times

🗳️ 👤 **GMartinelli** 3 years, 6 months ago

I don't know if its possible.. Since SGs are attached to the VPC, can you use it to reference on another VPC?

upvoted 1 times

🗳️ 👤 **2025flakyt** 3 years, 5 months ago

You can update security groups to reference peer VPC group

<https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-security-groups.html>