Amazon

Exam Questions AWS-Certified-Database-Specialty

AWS Certified Database - Specialty

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

     All examinations will be up to date.

* 24/7 Quality Support

     We will provide service round the clock.

* 100% Pass Rate

     Our guarantee that you will pass the exam.

* Unique Gurantee

     If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1
A business is transferring its on-premises database workloads to the Amazon Web Services (AWS) Cloud. A database professional migrating an Oracle database with a huge table to Amazon RDS has picked AWS DMS. The database professional observes that AWS DMS is consuming considerable time migrating the data. Which activities would increase the pace of data migration? (Select 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.

Answer: ACD

Explanation:
https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Tasks.LOBSupport.html

NEW QUESTION 2
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 ke
D. Apply the changes immediately.
E. Take a snapshot of the Aurora DB cluster and encrypt the snapshot using an AWS Key Management Service (AWS KMS) encryption ke
F. Restore the snapshot to a new DB cluster and update the financial application database endpoints.
G. Use AWS Key Management Service (AWS KMS) to secure the in-transit connection between the financial application and the Aurora DB cluster.

Answer: AB

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-replicas-adding.html

NEW QUESTION 3
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 group
B. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with sslmode=allow.
C. Set the rds.force_ssl=1 parameter in DB parameter group
D. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with sslmode=disable.
E. Set the rds.force_ssl=0 parameter in DB parameter group
F. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with sslmode=verify-ca.
G. Set the rds.force_ssl=1 parameter in DB parameter group
H. Download and use the Amazon RDS certificate bundle and configure the PostgreSQL connection string with sslmode=verify-full.

Answer: D

Explanation:
PostgreSQL: sslrootcert=rds-cert.pem sslmode=[verify-ca | verify-full]

NEW QUESTION 4
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

Answer: ACF

NEW QUESTION 5
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 replic
B. Reboot the read replic
C. Connect to the read replica and create the tables.

D. Change the read_only parameter to false (read_only=0) in the default parameter group of the read replic
E. Perform a reboot without failove
F. Connect to the read replica and create the tables using the local_only MySQL option.
G. Change the read_only parameter to false (read_only=0) in the default parameter grou
H. Reboot the read replic
I. Connect to the read replica and create the tables.
J. Create a new DB parameter grou
K. Change the read_only parameter to false (read_only=0). Associate the read replica with the new grou
L. Reboot the read replic
M. Connect to the read replica and create the tables.

Answer: D

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/rds-read-replica/

NEW QUESTION 6
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 30days.
D. Create an AWS Lambda function to run on the first day of every month to create an automated RDS snapshot.

Answer: A

Explanation:
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

NEW QUESTION 7
A business need a data warehouse system that stores data consistently and in a highly organized fashion. The organization demands rapid response times for end-user inquiries including current-year data, and users must have access to the whole 15-year dataset when necessary. Additionally, this solution must be able to manage a variable volume of incoming inquiries. Costs associated with storing the 100 TB of data must be maintained to a minimum.
Which solution satisfies these criteria?

A. Leverage an Amazon Redshift data warehouse solution using a dense storage instance type while keeping all the data on local Amazon Redshift storag
B. 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 dat
D. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum laye
E. Provision enough instances to support high demand.
F. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent dat
G. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum laye
H. Enable Amazon Redshift Concurrency Scaling.
I. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent dat
J. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum laye
K. Leverage Amazon Redshift elastic resize.

Answer: C

Explanation:
https://docs.aws.amazon.com/redshift/latest/dg/concurrency-scaling.html
"With 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."

NEW QUESTION 8
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 Serve
B. Use a CloudWatch dashboard to identify the root cause.
C. Enable RDS Performance Insights and determine which query is creating the proble
D. Request changes to the query to address the problem.
E. Use AWS X-Ray deployed with Amazon RDS to track query system traces.
F. Create a support request and work with AWS Support to identify the source of the issue.

Answer: B

Explanation:
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/

NEW QUESTION 9
A business maintains a SQL Server database on-premises. Active Directory authentication is used to provide users access to the database. The organization transferred their database successfully to Amazon RDS for SQL Server. The organization, however, has reservations regarding user authentication in the AWS Cloud environment.
Which authentication solution should a database professional provide?

A. Deploy Active Directory Federation Services (AD FS) on premises and configure it with an on-premises Active Director
B. Set up delegation between the on- premises AD FS and AWS Security Token Service (AWS STS) to map user identities to a role using theAmazonRDSDirectoryServiceAccess managed IAM policy.
C. Establish a forest trust between the on-premises Active Directory and AWS Directory Service for Microsoft Active Director
D. Use AWS SSO to configure an Active Directory user delegated to access the databases in RDS for SQL Server.
E. Use Active Directory Connector to redirect directory requests to the company€™s on-premises Active Directory without caching any information in the clou
F. Use the RDS master user credentials to connect to the DB instance and configure SQL Server logins and users from the Active Directory users and groups.
G. Establish a forest trust between the on-premises Active Directory and AWS Directory Service for Microsoft Active Director
H. Ensure RDS for SQL Server is using mixed mode authenticatio
I. Use the RDS master user credentials to connect to the DB instance and configure SQL Server logins and users from the Active Directory users and groups.

Answer: D

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

NEW QUESTION 10
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

Answer: C

Explanation:
https://docs.aws.amazon.com/documentdb/latest/developerguide/event-auditing.html Auditing Amazon DocumentDB Events
PDF
Kindle RSS
With Amazon DocumentDB (with MongoDB compatibility), you can audit events that were performed in your cluster. Examples of logged events include successful and failed authentication attempts, dropping a collection in a database, or creating an index. By default, auditing is disabled on Amazon DocumentDB and requires that you opt in to use this feature.
When auditing is enabled, Amazon DocumentDB records Data Definition Language (DDL), authentication, authorization, and user management events to Amazon CloudWatch Logs. 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.

NEW QUESTION 10
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 fil
B. Add the required corporate network IPs and remove the unwanted IPs.
C. Modify the associated security grou
D. Add the required corporate network IPs and remove the unwanted IPs.
E. Move the DB instance to a private subnet using AWS DMS.
F. Enable VPC peering between the application host running on the corporate network and the VPC associated with the DB instance.
G. Disable the publicly accessible setting.
H. Connect to the DB instance using private IPs and a VPN.

Answer: BEF

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.WorkingWithRDSInstanceinaVPC.ht

NEW QUESTION 13
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.

Answer: C

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

NEW QUESTION 16
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 instanc
B. Copy the encrypted snapshot to Amazon S3. Restore the DB instance from the encrypted snapshot using Amazon S3.
C. Create a new RDS for MySQL DB instance with encryption enable
D. Restore the unencrypted snapshot to this DB instance.
E. Create a snapshot of the unencrypted DB instanc
F. Create an encrypted copy of the snapsho
G. Restore the DB instance from the encrypted snapshot.
H. Temporarily shut down the unencrypted DB instanc
I. Enable AWS KMS encryption in the AWS Management Console using an AWS managed CM
J. Restart the DB instance in an encrypted state.

Answer: C

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html#Overview.Encryption.

NEW QUESTION 19
A significant automotive manufacturer is switching a mission-critical finance application's database to Amazon DynamoDB. According to the company's risk and compliance policy, any update to the database must be documented as a log entry for auditing purposes. Each minute, the system anticipates about 500,000 log entries. Log entries should be kept in Apache Parquet files in batches of at least 100,000 records per file.
How could a database professional approach these needs while using DynamoDB?

A. Enable Amazon DynamoDB Streams on the tabl
B. Create an AWS Lambda function triggered by the strea
C. Write the log entries to an Amazon S3 object.
D. Create a backup plan in AWS Backup to back up the DynamoDB table once a da
E. Create an AWS Lambda function that restores the backup in another table and compares both tables for change
F. Generate the log entries and write them to an Amazon S3 object.
G. Enable AWS CloudTrail logs on the tabl
H. Create an AWS Lambda function that reads the log files once an hour and filters DynamoDB API action
I. Write the filtered log files to Amazon S3.
J. Enable Amazon DynamoDB Streams on the tabl
K. Create an AWS Lambda function triggered by the strea
L. Write the log entries to an Amazon Kinesis Data Firehose delivery stream with buffering and Amazon S3 as the destination.

Answer: D

NEW QUESTION 20
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 AVS 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.

Answer: A

Explanation:
https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.Redshift.html

NEW QUESTION 22
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

Answer: D

Explanation:
https://docs.amazonaws.cn/en_us/AmazonRDS/latest/UserGuide/CHAP_BestPractices.html

NEW QUESTION 24
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.

Answer: B

Explanation:
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."

NEW QUESTION 26
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.

Answer: A

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Concepts.DBInstanceClass.html

NEW QUESTION 28
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 fil
E. Replace the passwords by running CloudFormation using a sed command.

Answer: B

Explanation:
https://aws.amazon.com/blogs/infrastructure-and-automation/securing-passwords-in-aws-quick-starts-using-aws

NEW QUESTION 29
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

Answer: B

NEW QUESTION 32
A small startup firm wishes to move a 4 TB MySQL database from on-premises to AWS through an Amazon RDS for MySQL DB instance.
Which migration approach would result in the LEAST amount of downtime?

A. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data cente
B. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucke
C. Import the snapshot into the DB instance utilizing the MySQL utilities running on an Amazon EC2 instanc
D. Immediately point the application to the DB instance.
E. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data cente
F. Use the mysqldump utility to create a snapshot of the on-premises MySQL serve
G. Copy the snapshot into the EC2 instance and restore it into the EC2 MySQL instanc

H. Use AWS DMS to migrate data into a new RDS for MySQL DB instanc
I. Point the application to the DB instance.
J. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data cente
K. Use the mysqldump utility to create a snapshot of the on-premises MySQL serve
L. 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 instanc
M. Point the application to the DB instance.
N. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data cente
O. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucke
P. Import the snapshot into the DB instance using the MySQL utilities running on an Amazon EC2 instanc
Q. Establish replication into the new DB instance using MySQL replicatio
R. Stop application access to the on-premises MySQL server and let the remaining transactions replicate ove
S. Point the application to the DB instance.

Answer: D

Explanation:
 https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/MySQL.Procedural.Importing.NonRDSRepl.html
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/MySQL.Procedural.Importing.External.Repl.html

NEW QUESTION 37
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

Answer: D

Explanation:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-partition-key-design.html#bp-partition "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."

NEW QUESTION 41
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.

Answer: A

Explanation:
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.

NEW QUESTION 46
A single MySQL database was moved to Amazon Aurora by a business. The production data is stored in a database cluster in VPC PROD, whereas 12 testing environments are hosted in VPC TEST with the same AWS account. Testing has a negligible effect on the test data. The development team requires that each environment be updated nightly to ensure that each test database has daily production data.
Which migration strategy will be the quickest and least expensive to implement?

A. Run the master in Amazon Aurora MySQ
B. Create 12 clones in VPC_TEST, and script the clones to be deleted and re-created nightly.
C. Run the master in Amazon Aurora MySQ
D. Take a nightly snapshot, and restore it into 12 databases in VPC_TEST using Aurora Serverless.
E. Run the master in Amazon Aurora MySQ
F. Create 12 Aurora Replicas in VPC_TEST, and script the replicas to be deleted and re-created nightly.
G. Run the master in Amazon Aurora MySQL using Aurora Serverles
H. Create 12 clones in VPC_TEST, and script the clones to be deleted and re-created nightly.

Answer: A

Explanation:

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

NEW QUESTION 48
Amazon Aurora MySQL is being used by an ecommerce business to migrate its main application database. The firm is now doing OLTP stress testing using concurrent database connections. A database professional detected sluggish performance for several particular write operations during the first round of testing. Examining the Amazon CloudWatch stats for the Aurora DB cluster revealed a CPU usage of 90%.
Which actions should the database professional take to determine the main cause of excessive CPU use and sluggish performance most effectively? (Select two.)

A. Enable Enhanced Monitoring at less than 30 seconds of granularity to review the operating system metrics before the next round of tests.
B. Review the VolumeBytesUsed metric in CloudWatch to see if there is a spike in write I/O.
C. Review Amazon RDS Performance Insights to identify the top SQL statements and wait events.
D. Review Amazon RDS API calls in AWS CloudTrail to identify long-running queries.
E. Enable Advance Auditing to log QUERY events in Amazon CloudWatch before the next round of tests.

Answer: AC

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/rds-instance-high-cpu/ https://aws.amazon.com/premiumsupport/knowledge-center/rds-mysql-slow-query/

NEW QUESTION 53
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

Answer: C

Explanation:
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

NEW QUESTION 56
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 instanc
B. 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.
C. 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 instanc
D. Configure an Amazon RDS event subscription to send a notification to an Amazon SNS topic to which the Lambda function is subscribed.
E. Configure one Aurora Replica to have the same instance class as the primary DB instanc
F. Implement Aurora PostgreSQL DB cluster cache managemen
G. Set the failover priority to tier-0 for the primary DB instance and one replica with the same instance clas
H. Set the failover priority to tier-1 for the otherreplicas.
I. Configure both Aurora Replicas to have the same instance class as the primary DB instanc
J. Implement Aurora PostgreSQL DB cluster cache managemen
K. Set the failover priority to tier-0 for the primary DB instance and to tier-1 for the replicas.

Answer: C

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.cluster-cache-mgmt.htm
https://aws.amazon.com/blogs/database/introduction-to-aurora-postgresql-cluster-cache-management/
"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. "
Amazon Aurora with PostgreSQL compatibility now supports cluster cache management, providing a faster path to full performance if there's a failover. With cluster cache management, you designate a specific reader DB instance in your Aurora PostgreSQL cluster as the failover target. Cluster cache management keeps the data in the designated reader's cache synchronized with the data in the read-write instance's cache. If a failover occurs, the designated reader is promoted to be the new read-write instance, and workloads benefit immediately from the data in its cache.

NEW QUESTION 59
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.

Answer: C

NEW QUESTION 64

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 reportin
B. Monitor and set up replication between the two clusters to keep data consistent.
C. Review all evaluate the peak combined workloa
D. Ensure that utilization of the DB cluster node is at an acceptable leve
E. Adjust the number of instances, if necessary.
F. Use the stop cluster functionality to stop all the nodes of the DB cluster during times of minimal workloa
G. The cluster can be restarted again depending on the workload at the time.
H. Set up automatic scaling on the DB cluste
I. This will allow the number of reader nodes to adjust automatically to the reporting workload, when needed.

Answer: D

NEW QUESTION 68

A business is transferring a database from one AWS Region to another using an Amazon RDS for SQL Server DB instance. The organization wishes to keep database downtime to a minimum throughout the transfer.

Which migration strategy should the organization use for this cross-regional move?

A. Back up the source database using native backup to an Amazon S3 bucket in the same Regio
B. Then restore the backup in the target Region.
C. Back up the source database using native backup to an Amazon S3 bucket in the same Regio
D. Use Amazon S3 Cross-Region Replication to copy the backup to an S3 bucket in the target Regio
E. Then restore the backup in the target Region.
F. Configure AWS Database Migration Service (AWS DMS) to replicate data between the source and the target database
G. Once the replication is in sync, terminate the DMS task.
H. Add an RDS for SQL Server cross-Region read replica in the target Regio
I. Once the replication is in sync, promote the read replica to master.

Answer: C

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.XRgn.html
With Amazon RDS, you can create a MariaDB, MySQL, Oracle, or PostgreSQL read replica in a different AWS Region from the source DB instance. Creating a cross-Region read replica isn't supported for SQL Server on Amazon RDS.

NEW QUESTION 72

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.

Answer: ABC

NEW QUESTION 74

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 tothe 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 connectio
D. Allow AWS DMS to handle syncing change data capture (CDC) data from the source to the target database.
E. 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 connectio
F. Allow AWS SCT to handle syncing change data capture (CDC) data from the source to the target database.

Answer: A

Explanation:
https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.S3.html Using Amazon S3 as a target for AWS Database Migration Service

NEW QUESTION 76
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 cente
B. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucke
C. Import the snapshot into the DB instance utilizing the MySQL utilities running on an Amazon EC2 instanc
D. Immediately point the application to the DB instance.
E. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data cente
F. Use the mysqldump utility to create a snapshot of the on-premises MySQL serve
G. Copy the snapshot into the EC2 instance and restore it into the EC2 MySQL instanc
H. Use AWS DMS to migrate data into a new RDS for MySQL DB instanc
I. Point the application to the DB instance.
J. Deploy a new Amazon EC2 instance, install the MySQL software on the EC2 instance, and configure networking for access from the on-premises data cente
K. Use the mysqldump utility to create a snapshot of the on-premises MySQL serve
L. 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 instanc
M. Point the application to the DB instance.
N. Deploy a new RDS for MySQL DB instance and configure it for access from the on-premises data cente
O. Use the mysqldump utility to create an initial snapshot from the on-premises MySQL server, and copy it to an Amazon S3 bucke
P. Import the snapshot into the DB instance using the MySQL utilities running on an Amazon EC2 instanc
Q. Establish replication into the new DB instance using MySQL replicatio
R. Stop application access to the on-premises MySQL server and let the remaining transactions replicate ove
S. Point the application to the DB instance.

Answer: B

NEW QUESTION 79
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 cluste
B. Launch an Amazon EC2 instance with a weekly cron job to execute the custom script.
C. 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 cluste
D. Schedule the Lambda function to run weekly using Amazon EventBridge (Amazon CloudWatch Events).
E. 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 cluste
F. Schedule both Lambda functions to run weekly using Amazon EventBridge (Amazon CloudWatch Events).
G. 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.

Answer: B

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Integrating.SaveIntoS3.htm

NEW QUESTION 83
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.

Answer: BD

NEW QUESTION 86
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.

Answer: C

Explanation:
"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

NEW QUESTION 89
Amazon DynamoDB global tables are being used by a business to power an online gaming game. The game is played by gamers from all around the globe. As the game became popularity, the amount of queries to DynamoDB substantially rose. Recently, gamers have complained about the game's condition being inconsistent between nations. A database professional notices that the ReplicationLatency metric for many replica tables is set to an abnormally high value.
Which strategy will resolve the issue?

A. Configure all replica tables to use DynamoDB auto scaling.
B. Configure a DynamoDB Accelerator (DAX) cluster on each of the replicas.
C. Configure the primary table to use DynamoDB auto scaling and the replica tables to use manually provisioned capacity.
D. Configure the table-level write throughput limit service quota to a higher value.

Answer: A

Explanation:
 https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/V2globaltables_reqs_bestpractices.html

NEW QUESTION 92
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

Answer: C

NEW QUESTION 94
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

Answer: BC

Explanation:
The following items need to be reconfigured after restoring the DynamoDB table.
--AutoScaling policy
--IAM policy
--CloudWatch settings
--Tags
--Stream settings
--TTL
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/backuprestore_HowItWorks.html

NEW QUESTION 95
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

Answer: B

Explanation:
 https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RebootInstance.html https://exain.wordpress.com/2017/07/12/amazon-rds-multi-az-setup-failover-simulation/
"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."

NEW QUESTION 100
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

Answer: A

NEW QUESTION 103
......

Relate Links

100% Pass Your AWS-Certified-Database-Specialty Exam with Exambible Prep Materials

https://www.exambible.com/AWS-Certified-Database-Specialty-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/